



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

**OAM AND DIAGNOSTICS GUIDE
RELEASE 14.0.R4**

3HE 10791 AAAB TQZZA 01

Issue: 01

July 2016

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.

© 2016 Nokia. All rights reserved.

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

Table of Contents

1	Getting Started	9
1.1	About This Guide	9
1.2	Router Configuration Process	10
2	Mirror Services	11
2.1	In This Chapter	11
2.2	Service Mirroring	12
2.3	Mirror Implementation.....	13
2.3.1	Mirror Source and Destinations	14
2.3.1.1	Local and Remote Mirroring	14
2.3.1.2	Slicing	15
2.3.2	Mirroring Performance	16
2.3.3	Mirroring Configuration	16
2.3.4	ATM Mirroring	18
2.3.5	IP Mirroring	19
2.3.5.1	Remote IP Mirroring	20
2.3.5.2	Local IP Mirroring	20
2.3.5.3	Port-ID Enabled PPP Mirroring.....	20
2.4	Mirrored Traffic Transport using MPLS-TP SDPs	21
2.5	Subscriber Mirroring	29
2.6	Lawful Intercept	30
2.6.1	LI Activation Through RADIUS	30
2.6.2	Routable Lawful Intercept Encapsulation	32
2.7	Pseudowire Redundant Mirror Services	36
2.7.1	Redundant Mirror Source Notes	38
2.8	Carrier Grade NAT – Lawful Intercept	39
2.9	Configuration Process Overview	40
2.10	Configuration Notes	41
2.11	Configuring Service Mirroring with CLI	43
2.11.1	Mirror Configuration Overview	43
2.11.1.1	Defining Mirrored Traffic	44
2.11.2	Lawful Intercept Configuration Overview	45
2.11.2.1	Saving LI Data	45
2.11.2.2	Regulating LI Access	46
2.11.2.3	Configurable Filter Lock for Lawful Intercept	50
2.11.2.4	LI MAC Filter Configuration	50
2.11.2.5	LI Logging	50
2.12	Basic Mirroring Configuration	51
2.12.1	Mirror Classification Rules	52
2.12.1.1	Port	53
2.12.1.2	SAP	54
2.12.1.3	MAC filter	55
2.12.1.4	IP filter	55
2.12.1.5	Ingress label	56
2.12.1.6	Subscriber	56

2.13	Common Configuration Tasks	57
2.13.1	Configuring a Local Mirror Service	58
2.13.2	Configuring SDPs for Mirrors and LI	59
2.13.3	Configuring a Remote Mirror Service	61
2.13.4	Configuring an ATM Mirror Service	63
2.13.5	Configuring Lawful Intercept Parameters	64
2.13.6	Pseudowire Redundancy for Mirror Services Configuration Example.....	65
2.14	Service Management Tasks	66
2.14.1	Modifying a Local Mirrored Service	67
2.14.2	Deleting a Local Mirrored Service	68
2.14.3	Modifying a Remote Mirrored Service	68
2.14.4	Deleting a Remote Mirrored Service	69
2.15	Mirror Service Configuration Command Reference.....	71
2.15.1	Command Hierarchies	71
2.15.1.1	Mirror Configuration Commands	71
2.15.1.2	Lawful Intercept Commands.....	73
2.15.2	Command Descriptions	76
2.15.2.1	Generic Commands.....	76
2.15.2.2	Mirror Destination Configuration Commands	78
2.15.2.3	Lawful Intercept Commands.....	101
2.16	Mirror Service Show and Debug Command Reference	141
2.16.1	Command Hierarchies.....	141
2.16.1.1	Show Commands	141
2.16.1.2	Debug Commands.....	141
2.16.2	Command Descriptions	142
2.16.2.1	Show Commands	142
2.16.2.2	Debug Commands.....	155
3	OAM, SAA, and OAM-PM	165
3.1	In This Chapter	165
3.2	OAM Overview	166
3.2.1	LSP Diagnostics: LSP Ping and Trace.....	167
3.2.2	LSP Ping/Trace for an LSP Using a BGP IPv4 Label Route	168
3.2.3	ECMP Considerations	169
3.2.4	Lsp-ping and Lsp-trace over Unnumbered IP Interface	172
3.2.5	Downstream Detailed Mapping (DDMAP) TLV	172
3.2.6	Using DDMAP TLV in LSP Stitching and LSP Hierarchy	175
3.2.6.1	Responder Node Procedures	176
3.2.6.2	Sender Node Procedures.....	177
3.2.7	LDP Tree Trace: End-to-End Testing of Paths in an LDP ECMP Network	179
3.2.8	LDP ECMP Tree Building	180
3.2.9	Periodic Path Exercising.....	181
3.2.10	LSP Ping for RSVP P2MP LSP (P2MP).....	181
3.2.11	LSP Trace for RSVP P2MP LSP	183
3.2.11.1	LSP Trace Behavior When S2L Path Traverses a Re-Merge Node.....	185
3.2.12	Tunneling of ICMP Reply Packets over MPLS LSP	187
3.2.13	QoS Handling of Tunneled ICMP Reply Packets	189

3.2.14	Summary of UDP Traceroute Behavior With and Without ICMP Tunneling	189
3.2.15	SDP Diagnostics.....	190
3.2.16	SDP Ping	191
3.2.17	SDP MTU Path Discovery	191
3.2.18	Service Diagnostics	192
3.2.19	VPLS MAC Diagnostics	192
3.2.20	MAC Ping	193
3.2.21	MAC Trace	193
3.2.22	CPE Ping	194
3.2.23	CPE Ping for PBB Epipе	195
3.2.23.1	Hardware Support	195
3.2.24	MAC Populate	196
3.2.25	MAC Purge	197
3.2.26	VLL Diagnostics.....	197
3.2.27	VCCV Ping	197
3.2.27.1	VCCV-Ping Application.....	197
3.2.27.2	VCCV Ping in a Multi-Segment Pseudowire.....	200
3.2.28	Automated VCCV-Trace Capability for MS-Pseudowire	201
3.2.28.1	VCCV for Static Pseudowire Segments	202
3.2.28.2	Detailed VCCV-Trace Operation	202
3.2.28.3	Control Plane Processing of a VCCV Echo Message in a MS-Pseudowire.....	203
3.2.29	IGMP Snooping Diagnostics.....	204
3.2.30	MFIB Ping.....	204
3.2.31	ATM Diagnostics	205
3.2.32	MPLS-TP On-Demand OAM Commands.....	206
3.2.33	MPLS-TP Pseudowires: VCCV-Ping/VCCV-Trace.....	206
3.2.33.1	VCCV Ping and VCCV Trace Between Static MPLS-TP and Dynamic PW Segments.....	207
3.2.34	MPLS-TP LSPs: LSP-Ping/LSP Trace	208
3.2.35	VxLAN Ping Supporting EVPN for VxLAN	210
3.2.36	Show Commands	210
3.2.37	BFD	210
3.3	IP Performance Monitoring (IP PM).....	213
3.3.1	Two-Way Active Measurement Protocol (TWAMP).....	213
3.3.2	Two-Way Active Measurement Protocol Light (TWAMP Light)	214
3.4	Ethernet Connectivity Fault Management (ETH-CFM).....	219
3.4.1	ETH-CFM Building Blocks	221
3.4.2	Loopback	238
3.4.3	Loopback Multicast.....	240
3.4.4	Linktrace	241
3.4.5	Continuity Check (CC)	243
3.4.6	CC Remote Peer Auto-Discovery	250
3.4.7	CCM Grace Period	252
3.4.8	CCM Hold Timers.....	253
3.4.9	Alarm Indication Signal (ETH-AIS Y.1731).....	254
3.4.10	Client Signal Fail (ETH-CSF Y.1731)	258
3.4.11	Test (ETH-TST Y.1731).....	259

3.4.12	One-Way Delay Measurement (ETH-1DM Y.1731)	260
3.4.13	Two-Way Delay Measurement (ETH-DMM Y.1731)	261
3.4.14	Synthetic Loss Measurement (ETH-SLM Y.1731).....	262
3.4.14.1	Configuration Example	264
3.4.15	Frame Loss Measurement (ETH-LMM Y.1731)	267
3.5	ETH-CFM Statistics	271
3.6	ETH-CFM Packet Debug	273
3.7	ETH-CFM CoS Considerations	275
3.8	OAM Mapping.....	276
3.8.1	CFM Connectivity Fault Conditions	277
3.8.2	CFM Fault Propagation Methods.....	278
3.8.3	Epipe Services.....	278
3.8.4	CFM Detected Fault	279
3.8.4.1	SAP/SDP-Binding Failure (Including Pseudowire Status)	279
3.8.4.2	Service Down	279
3.8.4.3	Interaction with Pseudowire Redundancy	280
3.8.5	Ipipe Services	280
3.8.5.1	CFM Detected Fault	280
3.8.5.2	SAP/SDP-binding Failure (Including Pseudowire Status)	305
3.8.5.3	Service Administratively Shutdown	305
3.8.5.4	Interaction with Pseudowire Redundancy	305
3.8.6	VPLS Service	305
3.8.6.1	CFM Detected Fault	305
3.8.6.2	SAP/SDP-Binding Failure (Including Pseudowire Status)	306
3.8.6.3	Service Down	306
3.8.6.4	Pseudowire Redundancy and Spanning Tree Protocol.....	307
3.8.7	IES and VPRN Services	307
3.8.8	Pseudowire Switching	307
3.8.9	LLF and CFM Fault Propagation	308
3.8.10	802.3ah EFM OAM Mapping and Interaction with Service Manager	308
3.9	Service Assurance Agent (SAA).....	309
3.10	OAM Performance Monitoring (OAM-PM).....	314
3.10.1	Session.....	316
3.10.2	Standard PM Packets.....	317
3.10.3	Detectable Transmit Errors.....	318
3.10.4	Measurement Intervals	320
3.10.5	Data Structures and Storage	330
3.10.6	Bin Groups.....	334
3.10.7	Relating the Components	336
3.10.8	IP Performance Monitoring	337
3.10.8.1	Accounting Policy Configuration	337
3.10.8.2	Service Configuration	337
3.10.8.3	OAM-PM Configuration	338
3.10.9	Ethernet Performance Monitoring	340
3.10.9.1	Accounting Policy Configuration	340
3.10.9.2	ETH-CFM Configuration	340
3.10.9.3	Service Configuration	340
3.10.9.4	Ethernet OAM-PM Configuration	341

3.10.10	OAM-PM Event Monitoring	344
3.11	Traceroute with ICMP Tunneling In Common Applications	350
3.11.1	BGP-LDP Stitching and ASBR/ABR/Data Path RR for BGP IPv4 Label Route	350
3.11.2	VPRN Inter-AS Option B	353
3.11.3	VPRN Inter-AS Option C and ASBR/ABR/Data Path RR for BGP IPv4 Label Route	355
3.12	Diagnostics Command Reference	359
3.12.1	Command Hierarchies	360
3.12.1.1	OAM Commands	360
3.12.1.2	SAA Commands	365
3.12.1.3	OAM Performance Monitoring and Binning Commands	367
3.12.1.4	IP Performance Monitoring Commands	370
3.12.1.5	Show Commands	373
3.12.1.6	Clear Commands	374
3.12.1.7	Monitor Commands	374
3.12.1.8	Debug Commands	374
3.12.1.9	Tools Commands	374
3.12.2	Command Descriptions	375
3.12.2.1	OAM and SAA Commands	375
3.12.2.2	Show Commands	545
3.12.2.3	Clear Commands	583
3.12.2.4	Monitor Commands	584
3.12.2.5	Debug Commands	589
3.12.2.6	Tools Commands	591
4	Standards and Protocol Support	593

1 Getting Started

1.1 About This Guide

This guide describes service mirroring and Operations, Administration and Management (OAM) and diagnostic tools provided by the router and presents examples to configure and implement various tests.

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.

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

- 7450 ESS
- 7750 SR
- 7950 XRS

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

Table 1 Supported SR OS Router Chassis Types

7450 ESS	7750 SR	7950 XRS
• 7450 ESS-6/6v • 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e	• 7950 XRS-16c • 7950 XRS-20/40

For a list of unsupported features by platform and chassis, refer to the *SR OS R14.0.Rx Software Release Notes*, part number 3HE10818 000x TQZZA.

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



Note: This guide generically covers Release 14.0 content and may contain some content that will be released in later maintenance loads. Please refer to the *SR OS R14.0.Rx Software Release Notes*, part number 3HE10818 000x TQZZA, for information on features supported in each load of the Release 14.0 software.

1.2 Router Configuration Process

[Table 2](#) lists the tasks necessary to configure mirroring, lawful intercept, 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 2 Configuration Process

Area	Task	Chapter
Diagnostics/ Service verification	Mirroring	Mirror Services
	Lawful Intercept	Lawful Intercept
	OAM	OAM, SAA, and OAM-PM
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support

2 Mirror Services

2.1 In This Chapter

This chapter provides information to configure mirroring.

Topics in this chapter include:

- [Service Mirroring](#)
- [Mirror Implementation](#)
 - [Mirror Source and Destinations](#)
 - [Local and Remote Mirroring](#)
 - [Slicing](#)
 - [Mirroring Performance](#)
 - [Mirroring Configuration](#)
 - [IP Mirroring](#)
- [Mirrored Traffic Transport using MPLS-TP SDPs](#)
- [Subscriber Mirroring](#)
- [Lawful Intercept](#)
- [Pseudowire Redundant Mirror Services](#)
- [Configuration Process Overview](#)
- [Configuration Notes](#)
- [Configuring Service Mirroring with CLI](#)
- [Basic Mirroring Configuration](#)
- [Common Configuration Tasks](#)
- [Service Management Tasks](#)
- [Mirror Service Configuration Command Reference](#)
- [Command Descriptions](#)

2.2 Service Mirroring

When troubleshooting complex operational problems, customer packets can be examined as they traverse the network. Nokia's service mirroring provides the capability to mirror customer packets to allow for trouble shooting and offline analysis. One way to accomplish this is with an overlay of network analyzers established at multiple PoPs, together with skilled technicians to operate them to decode the data provided. This method of traffic mirroring often requires setting up complex filters in multiple switches and/or routers. These, at best, are only able to mirror from one port to another on the same device.

Nokia's service mirroring extends and integrates these capabilities into the network and provides significant operational benefits. Each router can mirror packets from a specific service to any destination point in the network, regardless of interface type or speed.


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.

Nokia routers support service-based mirroring. While some Layer 3 switches and routers can mirror on a per-port basis within the device, Nokia routers can mirror on an n-to-1 unidirectional service basis and re-encapsulate the mirrored data for transport through the core network to another location, using either IP or MPLS tunneling as required ([Figure 1](#)).

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.

The mirrored frame size that is to be transmitted to the mirror destination can be explicitly configured by using slicing features. This enables mirroring only the parts needed for analysis. For example, only the headers can be copied for analysis, protecting the integrity and security of customer data, or conversely, copying the full packet, including customer data.

Figure 1 Service Mirroring



OSSG025

2.3 Mirror Implementation

Mirroring can be implemented on service access points (SAPs) or ingress network interfaces. The Flexible Fast Path processing complexes preserve the original packet throughout the forwarding and mirroring process, making any necessary packet changes, such as adding encapsulation, on a separate copy.

Nokia'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, the Flexible Fast Path network processor array (NPA) sends an exact copy of the original ingress packet 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 is forwarded to the mirror destination. Because the mirror copy of the packet is created before egress queuing, the mirrored packet stream may include copies of packets that are discarded in egress queues, such as during congestion or rate limiting.
- Mirroring supports tunnel destinations.

- 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.

2.3.1 Mirror Source and Destinations

Mirror sources and destinations have the following characteristics:

- They can be on the same SR OS (local) or on two different routers (remote).
- Mirror destinations can terminate on egress virtual ports which allows multiple mirror destinations to send to the same packet decode device, delimited by IEEE 802.1Q (referred to as Dot1q) tags. This is helpful when troubleshooting a multi-port issue within the network.

When multiple mirror destinations terminate on the same egress port, the individual dot1q tags can provide a DTE/DCE separation between the mirror sources.

- Packets ingressing a port can have a mirror destination separate from packets egressing another or the same port (the ports can be on separate nodes).
- Multiple mirror destinations are supported (local and/or remote) on a single chassis.
- The operational state of a mirror destination depends on the state of all the *outputs* of the mirror. The mirror destination will go operationally down if all the outputs are down (for example, all **mirror-dest>sap** and **mirror-dest>spoke-sdp** objects are down, and all gateways configured under **mirror-dest>encap** do not have a known route by which they can be reached). The state of a mirror destination does not depend on *inputs* such as SDPs configured under **mirror-dest>remote-source**, **debug>mirror-source** entries, or **config>li>li-source** entries. Some examples of outputs include **mirror-dest>sap** and **mirror-dest>spoke-sdp**.

2.3.1.1 Local and Remote Mirroring

Mirrored frames can be copied and sent to a specific local destination or service on the router (local mirroring) or copies can be encapsulated and sent to a different router (remote mirroring). This functionality allows network operators to centralize not only network analyzer (sniffer) resources, but also the technical staff who operate them.

The router allows multiple concurrent mirroring sessions so traffic from more than one ingress mirror source can be mirrored to the same or different egress mirror destinations.

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 wants to provide mirroring services. SDPs must be created first, before services can be configured.

2.3.1.2 Slicing

A further service mirroring refinement is “slicing” which copies a specified packet size of each frame. This is useful to monitor network usage without having to copy the actual data. Slicing enables mirroring larger frames than the destination packet decode equipment can handle. It also allows conservation of mirroring resources by limiting the size of the stream of packet through the router and the core network.

When a mirror **slice-size** is defined, a threshold that truncates a mirrored frame to a specific size is created. For example, if the value of 256 bytes is defined, up to the first 256 bytes of the frame are transmitted to the mirror destination. The original frame is not affected by the truncation. Mirrored frames, most likely, will grow larger as encapsulations are added when packets are transmitted through the network core or out the mirror destination SAP to the packet/protocol decode equipment. Note that slice-size is not supported by CEM encap-types or IP-mirroring (CEM encap-types applies to the 7750 SR and 7950 XRS only).

The transmission of a sliced or non-sliced frame is also dependent on the mirror destination SDP path MTU and/or the mirror destination SAP physical MTU. Packets that require a larger MTU than the mirroring destination supports are discarded if the defined slice size does not truncate the packet to an acceptable size.

2.3.2 Mirroring Performance

Replication of mirrored packets can, typically, affect performance and should be used carefully. Nokia routers minimize the impact of mirroring on performance by taking advantage of its distributed Flexible Fast Path technology. Flexible Fast Path forwarding allows efficient mirror service scaling and, at the same time, allows a large amount of data to be mirrored with minimal performance impact. When a mirror destination is configured, the packet slice option can truncate mirrored packets to the destination, which minimizes replication and tunneling overhead.

2.3.3 Mirroring Configuration

Mirroring can be performed based on the following criteria:

- Port
- SAP
- MAC filter
- IP filter
- Ingress label
- Subscriber


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 shows a local mirror service configured on ALA-A.

- Port 2/1/2 is specified as the source. Mirrored traffic ingressing and egressing this port will be sent to port 2/1/3.
- SAP 2/1/3 is specified as the destination. The sniffer is physically connected to this port. Mirrored traffic ingressing and egressing port 2/1/2 is sent here. SAP, encapsulation requirements, packet slicing, and mirror classification parameters are configured. SDPs are not used in local mirroring.

Figure 2 Local Mirroring Example



OSSG026

Figure 3 shows a remote mirror service configured as ALA B as the mirror source and ALA A as the mirror destination. Mirrored traffic ingressing 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 ingressing and egressing this port.

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

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



OSSG027

2.3.4 ATM Mirroring

ATM mirror functionality allows 7750 SR users to mirror AAL5 packets from a source ATM SAP to a destination ATM SAP connected locally or remotely. This functionality can be used to monitor the ATM traffic on a particular ATM SAP. In both the local and remote scenarios the source and destination SAPs must be of ATM SAP type.

All ingress and egress AAL5 traffic at the source ATM SAP is duplicated and sent toward the destination ATM SAP. Mirroring the ingress traffic only, egress traffic only, or both, can be configured. ATM OAM traffic is not mirrored toward the destination ATM SAP.

IP filters used as a mirror source are supported on ATM SAPs based on the IP filter applicability for different services.

ATM mirroring is applicable to the following services using an ATM SAP:

- Layer 3: IES and VPRN
- Layer 2: Apipe (sdu-type only), lpipe, Epipe, VPLS

ATM mirroring on an ATM SAP extends the service mirroring feature to include mirror sources with SAP type of ATM. Mirroring is supported on the following services:

- IES
- VPRN
- VPLS
- Epipe
- lpipe
- Apipe VLL service with the AAL5 SDU mode (atm-sdu spoke-sdp type)

Characteristics include:

- Supported only ATM MDAs and on the Any Service Any Port (ASAP) MDA.
- Mirror destinations for ATM mirroring must be ATM SAPs and cannot be part of an APS group, an IMA bundle, or an IMA Bundle Protection Group (BPGRP).
- A mirror source can be an ATM SAP component of an IMA bundle but cannot be part of an IMA BPGRP.
- ATM SAPs of an Apipe service with N:1 cell mode (atm-vcc, atm-vpc, and atm-cell spoke-sdp types) cannot be ATM mirror sources.

In [Figure 4](#), CE 3 is connected to PE1 on ATM SAP 2/1/1:0/100 as part of an IES service. The traffic on ATM SAP 2/1/1:0/100 is mirrored locally to CE4 device through ATM SAP 1/2/1:1/101. In this scenario, all AAL5 packets arriving at SAP 2/1/1:0/100 are duplicated and sent towards ATM SAP 1/2/1:1/101.

Figure 4 Example of an ATM Mirror Service




Fig_21

In the case where the destination ATM SAP is on a remote node PE2, then the AAL5 traffic arriving at ATM SAP 2/1/1:0/100 is duplicated and sent across the IP/MPLS network to PE2. At PE2 the traffic is forwarded to ATM SAP 1/1/1:0/1000 towards the ATM traffic monitoring device.

2.3.5 IP Mirroring

The IP mirroring capability for the 7750 SR and 7950 XRS allows a mirror to be created with a parameter that specifies that only the IP packet is mirrored without the original ATM/FR/POS/Ethernet DLC header. This results in the mirrored IP packet becoming media agnostic on the mirror service egress.

This option is configurable on SAP mirrors for IES, VPRN and VPLS services, lpipe services, and subscriber mirrors. It is not supported on VLL services such as Apipe, Epipe, Fpipe, and on ports.

2.3.5.1 Remote IP Mirroring

With remote IP mirroring, the mirror destination configuration can allow IP packets to be mirrored from a source router (Figure 5). The packets will be delivered to the destination in a spoke-terminated interface created in a VPRN service. IES interfaces are not supported. The interface can be configured with policy-based routing filters to allow sniffer selection based on incoming mirrored destination IP addresses. The interface cannot send traffic out as it is a destination only feature. Packets arriving at the interface will be routed based on the routing information within the VPRN. Policy-based routing should always be used unless only a sniffer is connected to the VPRN.

Figure 5 Remote IP Mirroring




Fig_17

2.3.5.2 Local IP Mirroring

Local mirroring is similar to remote mirroring but the source and destination of the mirror exist in the same Local IP mirroring node. The configuration must include the source address and destination MAC addresses for the packets going to the sniffer. The destination SAP must be Ethernet.

2.3.5.3 Port-ID Enabled PPP Mirroring

Operators that use mirroring for statistics collection make use of VLANs or DLCIs for customer separation. Since PPP offers no such separation, the maximum number of PPP circuits may be identified (one per destination). This feature provides a proprietary mechanism to allow a single mirror to be used and only applies to the 7450 ESS and 7750 SR.

Port-ID enabled PPP mirroring includes the system's port ID in the mirrored packet. An operator using this flag in a PPP mirror will be able to identify the end customer circuit by finding the system's port ID (which is optionally made persistent) and correlating it to the port-id in the mirrored packet.

This mirroring does not change the priority of the mirror order (port/sap/sub/filter). Lawful intercept mirrors can use the flag and their priority is also maintained.

Since the inclusion of the port ID flag is placed on the mirror destination, all mirrored packets of all sources will include the port ID. For remote mirroring, the mirror destination service at the source node must be configured with this flag.

Note the following restrictions:

- This flag can only be used with a PPP mirror destination.
- This flag is mutually exclusive with a remote-source.
- This flag cannot be enabled on a an IP mirror type.


2.4 Mirrored Traffic Transport using MPLS-TP SDPs

Bidirectional MPLS-TP spoke SDPs with a configured pw-path-id can transport a mirrored service. Mirror services are not supported on static PWs with an MPLS-TP pw-path-id bound to an SDP that uses an RSVP-TE LSP.

Mirror services using MPLS-TP spoke SDPs can be configured using CLI in the context mirror-dest>remote-source. For both the CPM and IOM, this enables reuse of spokes for mirror services and other services such as pipes.

Control channel status signaling is supported with PW redundancy on spoke SDPs in a mirror context.

The following is an example of PW redundancy for a mirror service. In this case, MPLS-TP spoke SDPs are used.

Figure 6 Mirroring with PW Redundancy using MPLS-TP

al_0526

Note that mirroring traffic is usually unidirectional, flowing from “source” nodes (B or C) to “destination” nodes (D or E). However in case of MPLS-TP, the control channel status packets may flow in the reverse direction.

The following is an example of a mirror service configuration using MPLS-TP spoke SDPs:

Source Node B

```
#-----
echo "Mirror Configuration"
#-----

mirror
  mirror-dest 300 create
  endpoint "X" create
  revert-time 100
  exit
  endpoint "Y" create
  revert-time 100
  exit
  remote-source
    spoke-sdp 230:1300 endpoint "Y" icb create
    ingress
      vc-label 13301
    exit
    egress
      vc-label 13301
    exit
    control-word
    pw-path-id
      agi 1:1
      saii-type2 1:10.20.1.2:13301
      taii-type2 1:10.20.1.3:13301
    exit
    control-channel-status
```

```
        refresh-timer 10
        no shutdown
    exit
    no shutdown
exit
spoke-sdp 240:300 endpoint "X" create
    ingress
        vc-label 2401
    exit
    egress
        vc-label 2401
    exit
    control-word
    pw-path-id
        agi 1:1
        saii-type2 1:10.20.1.2:2401
        taii-type2 1:10.20.1.4:2401
    exit
    control-channel-status
        refresh-timer 10
        no shutdown
    exit
    no shutdown
exit
spoke-sdp 250:300 endpoint "X" create
    ingress
        vc-label 6501
    exit
    egress
        vc-label 6501
    exit
    control-word
    pw-path-id
        agi 1:1
        saii-type2 1:10.20.1.2:6501
        taii-type2 1:10.20.1.5:6501
    exit
    control-channel-status
        refresh-timer 10
        no shutdown
    exit
    no shutdown
exit
spoke-sdp 230:300 endpoint "X" icb create
    ingress
        vc-label 12301
    exit
    egress
        vc-label 12301
    exit
    control-word
    pw-path-id
        agi 1:1
        saii-type2 1:10.20.1.2:12301
        taii-type2 1:10.20.1.3:12301
    exit
    control-channel-status
        refresh-timer 10
```

```

        no shutdown
        exit
        no shutdown
        exit
        no shutdown
        exit
        exit
    exit all

```

Destination Node C

```

#-----
echo "Mirror Configuration"
#-----
    mirror
        mirror-dest 300 create
            endpoint "X" create
                revert-time 100
            exit
            endpoint "Y" create
                revert-time 100
            exit
        remote-source
            spoke-sdp 230:1300 endpoint "Y" ich create
                ingress
                    vc-label 13301
                exit
                egress
                    vc-label 13301
                exit
                control-word
                pw-path-id
                    agi 1:1
                    saii-type2 1:10.20.1.3:13301
                    taii-type2 1:10.20.1.2:13301
                exit
                control-channel-status
                    refresh-timer 10
                    no shutdown
                exit
                no shutdown
            exit
        exit
    spoke-sdp 340:300 endpoint "X" create
        ingress
            vc-label 6501
        exit
        egress
            vc-label 6501
        exit
        control-word
        pw-path-id
            agi 1:1
            saii-type2 1:10.20.1.3:6501
            taii-type2 1:10.20.1.4:6501
        exit
        control-channel-status
            refresh-timer 10

```



```

        no shutdown
    exit
    no shutdown
exit
spoke-sdp 350:300 endpoint "X" create
    ingress
        vc-label 2401
    exit
    egress
        vc-label 2401
    exit
    control-word
    pw-path-id
        agi 1:1
        saii-type2 1:10.20.1.3:2401
        taii-type2 1:10.20.1.5:2401
    exit
    control-channel-status
        refresh-timer 10
        no shutdown
    exit
    no shutdown
exit
spoke-sdp 230:300 endpoint "X" icb create
    ingress
        vc-label 12301
    exit
    egress
        vc-label 12301
    exit
    control-word
    pw-path-id
        agi 1:1
        saii-type2 1:10.20.1.3:12301
        taii-type2 1:10.20.1.2:12301
    exit
    control-channel-status
        refresh-timer 10
        no shutdown
    exit
    no shutdown
exit
    no shutdown
exit
exit

```

Source Node D

```

#-----
echo "Mirror Configuration"
#-----
    mirror
        mirror-dest 300 create
        endpoint "X" create
        revert-time 100
    exit
    endpoint "Y" create
        revert-time 100

```

```
exit
remote-source
spoke-sdp 240:300 endpoint "Y" create
  ingress
    vc-label 2401
  exit
  egress
    vc-label 2401
  exit
  control-word
  pw-path-id
    agi 1:1
    saii-type2 1:10.20.1.4:2401
    taii-type2 1:10.20.1.2:2401
  exit
  control-channel-status
    refresh-timer 10
    no shutdown
  exit
  no shutdown
exit
spoke-sdp 340:300 endpoint "Y" create
  ingress
    vc-label 6501
  exit
  egress
    vc-label 6501
  exit
  control-word
  pw-path-id
    agi 1:1
    saii-type2 1:10.20.1.4:6501
    taii-type2 1:10.20.1.3:6501
  exit
  control-channel-status
    refresh-timer 10
    no shutdown
  exit
  no shutdown
exit
spoke-sdp 450:1300 endpoint "Y" icb create
  ingress
    vc-label 13301
  exit
  egress
    vc-label 13301
  exit
  control-word
  pw-path-id
    agi 1:1
    saii-type2 1:10.20.1.4:13301
    taii-type2 1:10.20.1.5:13301
  exit
  control-channel-status
    refresh-timer 10
    no shutdown
  exit
  no shutdown
exit
```

```

exit
sap lag-10:300.1 endpoint "X" create
exit
spoke-sdp 450:300 endpoint "X" icb create
  ingress
    vc-label 12301
  exit
  egress
    vc-label 12301
  exit
  control-word
  pw-path-id
    agi 1:1
    saii-type2 1:10.20.1.4:12301
    taii-type2 1:10.20.1.5:12301
  exit
  control-channel-status
    refresh-timer 10
    no shutdown
  exit
  no shutdown
exit
no shutdown
exit
exit

```

Destination Node E

```

#-----
echo "Mirror Configuration"
#-----
  mirror
    mirror-dest 300 create
    endpoint "X" create
    revert-time 100
  exit
  endpoint "Y" create
    revert-time 100
  exit
  remote-source
    spoke-sdp 250:300 endpoint "Y" create
    ingress
      vc-label 6501
    exit
    egress
      vc-label 6501
    exit
    control-word
    pw-path-id
      agi 1:1
      saii-type2 1:10.20.1.5:6501
      taii-type2 1:10.20.1.2:6501
    exit
    control-channel-status
      refresh-timer 10
      no shutdown
    exit
    no shutdown

```

```
exit
spoke-sdp 350:300 endpoint "Y" create
    ingress
        vc-label 2401
    exit
    egress
        vc-label 2401
    exit
    control-word
    pw-path-id
        agi 1:1
        saii-type2 1:10.20.1.5:2401
        taii-type2 1:10.20.1.3:2401
    exit
    control-channel-status
        refresh-timer 10
        no shutdown
    exit
    no shutdown
exit
spoke-sdp 450:1300 endpoint "Y" icb create
    ingress
        vc-label 13301
    exit
    egress
        vc-label 13301
    exit
    control-word
    pw-path-id
        agi 1:1
        saii-type2 1:10.20.1.5:13301
        taii-type2 1:10.20.1.4:13301
    exit
    control-channel-status
        refresh-timer 10
        no shutdown
    exit
    no shutdown
exit
exit
sap lag-10:300.1 endpoint "X" create
exit
spoke-sdp 450:300 endpoint "X" icb create
    ingress
        vc-label 12301
    exit
    egress
        vc-label 12301
    exit
    control-word
    pw-path-id
        agi 1:1
        saii-type2 1:10.20.1.5:12301
        taii-type2 1:10.20.1.4:12301
    exit
    control-channel-status
        refresh-timer 10
        no shutdown
    exit
```

```
        no shutdown
    exit
        no shutdown
    exit
exit
```

2.5 Subscriber Mirroring

This section describes mirroring based on a subscriber match. Subscriber mirroring applies only to the 7450 ESS and 7750 SR. Enhanced subscriber management provides the mechanism to associate subscriber hosts with queuing and filtering resources in a shared SAP environment. Mirroring used in subscriber aggregation networks for lawful intercept and debugging is required. With this feature, the mirroring capability allows the match criteria to include a subscriber-id.

Subscriber mirroring provides the ability to create a mirror source with subscriber information as match criteria. Specific subscriber packets can be mirrored mirror when using ESM with a shared SAP without prior knowledge of their IP or MAC addresses and without concern that they may change. The subscriber mirroring decision is more specific than a SAP. If a SAP (or port) is placed in a mirror and a subscriber host of which a mirror was configured is mirrored on that SAP packets matching the subscriber host will be mirrored to the subscriber mirror destination.

The mirroring configuration can be limited to specific forwarding classes used by the subscriber. When a forwarding class (FC) map is placed on the mirror only packets that match the specified FCs are mirrored. A subscriber can be referenced in maximum 2 different mirror-destinations: 1 for ingress and 1 for egress.

Subscriber based criteria in a mirror source remains in the mirror/li source configuration even if the subscriber is deleted, removed or logs off. When the subscriber returns (is configured/created or logs in) the mirroring will resume. This also implies that a subscriber can be configured as a mirror/li source before the actual subscriber exists on the node and before the subscriber id is active (the mirroring will start once the subscriber is actually created or logs in and the subscriber id becomes active).

2.6 Lawful Intercept

Lawful Intercept (LI) describes a process to intercept telecommunications by which law enforcement authorities can un-obtrusively monitor voice and data communications to combat crime and terrorism with higher security standards of lawful intercept capabilities in accordance with local law and after following due process and receiving proper authorization from competent authorities. The interception capabilities are sought by various telecommunications providers.

As lawful interception is subject to national regulation, requirements vary from one country to another. Nokia's implementation satisfies most national standard's requirements. LI capability is configurable for all Nokia service types.

LI mirroring is configured by an operator that has LI permission. LI mirroring is hidden from anyone who does not have the right permission.

2.6.1 LI Activation Through RADIUS

In addition to CLI and SNMP control, RADIUS messages also activate LI sessions for subscriber-host targets. Activation through RADIUS is equivalent to adding or removing a set of subscriber-host entries in an li-source.



Note: The term “activation” in this section represents both “activation and de-activation”.

The activation of an LI session via RADIUS applies to the 7450 ESS and 7750 SR and can occur in one of two ways:

- At the time the RADIUS access-accept message is received by the 7450 ESS or 7750 SR. In this case, the target (either a host, or a set of hosts) is implicit. The target acts as the same host (or set of hosts) that is within the scope of the access-accept and interception occurs for this entire set of hosts (or single host).
- “Through RADIUS COA messages. In this case, the target (set of hosts) is identified through one of the following methods:
 - acct-session-id (which can represent a single host or a collection of hosts), or
 - a <sap-id;ip-addr> carried in the NAS-Port-Id (attr 87) and the Framed-Ip-Address (attr 8).” for IPv4 hosts, or

- a <sap-id;IPv6_addr> carried in the NAS-Port-ID (attr 87) and one of Alc-Ipv6-Address, Framed-Ipv6-Prefix, or Delegated-Ipv6-Prefix for IPv6 hosts, or
- alc-subsc-id-str

The following set of VSAs are used to activate LI sessions via RADIUS:

- ALC-LI-Action – ON/OFF/NONE
- ALC-LI-Dest - <string>
 - The number is in ASCII format indicating mirror service
 - Future development will extend the definition of the handle to be attached to intercepted packets of the given subscriber-host
- ALC-LI-Direction – INGRESS/EGRESS
- ALC-LI-FC – be/l1/l2/af/ef

The ALC-LI-FC-MAP VSA can be present several times if more than one forwarding class (FC) is subject to LI.

ALC-LI-Direction and ALC-LI-FC are optional. If either is not included, both directions (ingress and egress) as well as all FCs will be mirrored.

Including the above VSAs in access-accept message will activate LI for newly created host. Note that in this case, the LI activation is not addressed by acct-session-id as this is not yet known during session authorization.

Different attributes can be used in a CoA to identify one or multiple subscriber host(s). Typically only a single (set of) attribute(s) is used to target a host or a number of hosts: "NAS-Port-Id + IP" or "Acct-Session-Id" or "Alc-Subsc-ID-Str". In case that both "NAS-Port-Id + IP" is used in a Wholesale/Retail model then the Alc-Retail-Serv-Id VSA must be included in the CoA.

The ability to delete all li-source entries from a particular mirror service is also available via RADIUS. This function may be useful when an LI Mediation device loses sync with the state of the SR OS and needs to reset a mirror service to a known state with no LI sessions. This clear function is performed by sending the following attributes in a RADIUS CoA. If the CoA does not contain exactly these three VSAs (each with a valid value matching the configuration on the SR OS) then the CoA will be silently dropped with no NAK:

- ALC-LI-Action = clear-dest-service
- ALC-LI-Dest-Service = *service-id*
- ALC-Authentication-Policy-Name (the authentication policy name used to authenticate the subscribers)

The LI-related VSA cannot be combined in one CoA message with other action-related VSAs (force-renew, change of sla-profile, etc.). The only exception to this rule is for the CoA used to create new sub-host. Then, LI-related VSAs can be included along with other VSAs.


If LI is activated through CLI/SNMP, the activation through RADIUS takes precedence. The precedence in this context means that RADIUS activation of LI will fully override whatever was configured at CLI/SNMP level for this particular host. If the RADIUS LI is de-activated, the CLI/SNMP configuration will become active again.

The LI-related VSAs are not shown in debug messages. The **show li li-source <xxxx>** command shows all sub-hosts for which LI was activated using RADIUS VSAs. This command is only accessible to CLI users with LI privileges.

2.6.2 Routable Lawful Intercept Encapsulation


The Routable LI encapsulation feature allows LI mirrored packets to be placed into a routable (for example, IP/UDP) header and then forwarded in a routing context (base or VPRN). An LI-shim inserted before the customer packet allows correlation of packets to LI sessions at the downstream LI Mediation device (LIG).

Figure 7 Routable Lawful Intercept Encapsulation



OSSG687


Figure 8 Routable Encapsulation Format



OSSG685

Some of the supported attributes and scenarios for the routable LI encapsulation feature include the following:

- The part of the customer packet that is copied and placed into the routable encapsulation can be either the IP packet (with none of the original Layer2 encap) or an Ethernet packet by selecting either ip-only or ether as the mirror-dest type.
- The ability to inject into the Base routing instance (for forwarding out network interfaces or IES SAPs for example) or a VPRN service.
- The ability to forward the encapsulated packets out VPRN SDPs, IGP/BGP shortcuts and SDP spoke interfaces.
- Options to use ip, udp, li-shim or ip, gre routable encapsulation (applies to the 7450 ESS and 7750 SR).
- An optional direction bit in the li-shim.
 - If the use of the direction bit is configured, then a bit from the intercept-id (config under the mirror-dest) is “stolen”. Only a 29b intercept-id will be allowed for li-source entries if the mirror-dest is configured to use a direction-bit.

Figure 9 LI-Shim version 01 with a direction bit

OSSG686

- The encoding of the direction (d) bit is as follows:
 - 0 = ingress
 - 1 = egress
- For NAT based LI, ingress means the traffic arriving at the node from the subscriber host (applies to the 7450 ESS and 7750 SR).
- User configurable **intercept-id** and **session-id** per li-source entry that is placed into the li-shim (a total max of 62 configurable bits).
- Configuration via CLI/SNMP or RADIUS (applies to the 7450 ESS and 7750 SR). For RADIUS configuration the following VSAs are used:
 - Alc-LI-Action, Alc-LI-Direction, Alc-LI-Destination, Alc-LI-FC: See the section called “LI Activation Through RADIUS” in this document for details.
 - Alc-LI-Intercept-Id: specifies the intercept-id to place in the LI-Shim. Only applicable if the mirror-dest (as specified by the Alc-LI-Destination) is configured with routable encap that contains the LI-Shim. A value of 0 is used if this VSA is not present.
 - Alc-LI-Session-Id: specifies the session-id to place in the LI-Shim. Only applicable if the mirror-dest (as specified by the Alc-LI-Destination) is configured with routable encap that contains the LI-Shim. A value of 0 is used if this VSA is not present.
- A LI session configured via RADIUS takes precedence over a session configured via CLI, but the CLI mirror is re-instated if the RADIUS mirror request is later removed (applies to the 7450 ESS and 7750 SR)
- ip | udp | li-shim encap is available for ether and ip-only mirror-dest types (note that ip-only supports, amongst other formats, packets that are reassembled from ATM cells.)
- ip | udp | li-shim encap is available for all li-source entry types: sap, filter, subscriber and nat.
 - Note that for NAT based Lawful Intercept, routable LI encap is available, as well as the mac/l2 based encap for NAT LI as configured under **config>li>li-source>nat>ethernet-encap** (applies to the 7450 ESS and 7750 SR)

- Fragmentation of the resulting mirror packet is supported. Note that fragmentation is supported for NAT LI with the Routable Encapsulation, but fragmentation is not supported for NAT LI with ethernet-encap (applies to the 7450 ESS and 7750 SR).

The following restrictions apply to the routable LI encapsulation feature:

- Only applicable to Lawful Intercept and is not available for debug or MS-ISA based Application Assurance mirrors. MS-ISA based Application Assurance is applicable to the 7450 ESS and 7750 SR.
- Not applicable to frame-relay, PPP, ATM-SDU, SAToP, or CESoPSN mirror-dest types.
- IPv4 transport only (the routable encapsulation cannot be IPv6).
- On the mirror source node, forwarding of routable encapsulated LI packets out of an R-VPLS interface is not supported. A mirror-dest configured with routable encapsulation can be bound to a routing instance that also has an R-VPLS bound to it, but the operator must ensure that the destination of the LI packets is not reachable via any R-VPLS interfaces. Any routable encapsulated LI packets that arrive at the egress of an R-VPLS interface are discarded. Parallel use of routable LI encapsulation and R-VPLS in the same routing instance is supported as long as the mirrored packets don't egress out of the R-VPLS interface.
- ip | gre encap is supported for the ip-only mirror-dest type only, and only for subscriber li-source entries (CLI/SNMP or RADIUS based). Subscriber management is not supported on the 7950 XRS.
 - The contents of the GRE header are all zeros (all optional bits zero, no optional headers/fields like checksum, offset, key, seq, etc.) except for the Protocol field which will contain 0x0800 for IPv4 packets or 0x86DD for IPv6 packets. The far end receiver of the intercepted packets must be configured to expect no GRE options (that is, no key, no checksum, etc.).
- On the mirror source 7450 ESS or 7750 SR node, both the card where the mirroring occurs and the card where the mirrored packet egresses the node must be FP2 based (IOM3-XP or IMMs). When employing a filter-based LI, Chassis Mode D must be used. For all other types of LI it is strongly recommended to use Chassis Mode D with this feature. If Chassis Mode D is not possible then extreme care by the operator will need to be employed to ensure that all possible interceptions can only occur on FP2 based cards, and that all possible outgoing interfaces for the mirrored/encapsulated packets are on FP2 based cards.

- On the source node where LI mirroring occurs, the operator must configure the mirror-dest to inject into the routing instance (that is, base or VPRN) in which the actual destination address is reachable *without* having to hop into a different instance using GRT leaking. In other words the interface out which the packet will end up traveling must exist in the routing instance that is configured in the mirror-dest.
 - For example, if the LIG is at 110.120.130.140 and is in the base instance, but VPRN-1 has a default route to the GRT (for example, 0.0.0.0->GRT) then the operator must configure the mirror-dest to inject into the base (even though theoretically address 110.120.130.140 is reachable from VPRN-1). If the operator attempts to configure the mirror-dest to inject into VPRN-1, and VPRN-1 itself does not have reachability to 110.120.130.140 out an interface that is part of the VPRN, then the mirror-dest will be operationally down.
- Platforms: Not supported on the 7450 ESS-1.


Care must be taken in the configuration of LI mirrors and the destination IP address for the routable LI encapsulation. Incorrect selection of the destination IP could send packets to unintended destinations (for example - configuring the encapsulation with a subscriber's IP address), and combinations of mirrors and routable encapsulation can create loops in the network.

2.7 Pseudowire Redundant Mirror Services

This section describes the implementation and configuration of redundant Mirror/Lawful Intercept services using redundant pseudowires.

Regardless of the protection mechanism (MC-LAG, STP or APS) the source switch will only transmit on the active link and not simultaneously on the standby link. As a result when configuring a redundant mirror / LI service or a mirror service where the customer has a redundant service but the mirror / LI service is not redundant the mirror source must be configured on both (A and B) PE nodes. In either case the PE with a mirror source will establish a pseudo wire to each eligible PE where the mirror / LI service terminates.


Figure 10 State Engine for Redundant Service to a Redundant Mirror Service



It is important to note that in order to provide protection in case the active SDP between node A and D fails and the need to limit the number of lost data for LI the ICB between node A and B must be supported. As a result when the SDP connecting nodes A and D fails the data on its way from the source switch to node A and the data in node A must be directed by the ICB to node B and from there to node D.


This functionality is already supported in when providing pseudo wire redundancy for VLLs and must be extended to mirror / LI service redundancy.

Figure 11 State Engine for Redundant Service to a Non-Redundant Mirror Service



The notable difference with scenarios standard pseudo wire redundancy scenarios is that provided the customer service is redundant on nodes A and B ([Figure 10](#) and [Figure 11](#)) both aggregation node A and Aggregation node B maintain an active Pseudo wire to Node D who in turn has an active link to the destination switch. If in the sample in [Figure 10](#), the link between D and the destination switch is disconnected then both aggregation A and B must switch to use pseudo wire connection to Node C.

Figure 12 State Engine for a Non-Redundant Service to a Redundant Mirror Service



In the case where a non redundant service is being mirrored to a redundant mirror service ([Figure 12](#)) the source aggregation node (A) can only maintain a pseudo wire to the active destination aggregation node (D). Should the link between aggregation node D and the destination switch fail then the pseudo wire must switch to the new active aggregation node (C).

2.7.1 Redundant Mirror Source Notes

A redundant remote mirror service destination is not supported for IP Mirrors (a set of remote IP mirror destinations). The remote destination of an IP mirror is a VPRN instance, and an “endpoint” cannot be configured in a VPRN service.

A redundant mirror source is supported for IP mirrors, but the remote destination must be a single node (a set of mirror source nodes, each with a mirror destination that points to the same destination node). In this case the destination node would have a VPRN instance with multiple ip-mirror-interfaces.

Multi Chassis APS (MC-APS) groups can not be used as the SAP for a redundant remote mirror destination service. APS can not be used to connect the remote mirror destination SR nodes to a destination switch.

Multi Chassis APS (MC-APS) groups can be used as the SAP for a redundant mirror service source. APS can be used to redundantly connect the source of the mirrored traffic to the SR nodes that are behaving as the mirror-sources.

2.8 Carrier Grade NAT – Lawful Intercept

Lawful intercept for NAT is supported to mirror configured subscriber’s traffic to a mirror-destination. When active, packets are mirrored from the perspective of the NAT outside interface (thus after NAT translations have occurred). All traffic for the specified subscriber, including traffic associated with static port-forwards, is mirrored. This feature is supported for the 7450 ESS and 7750 SR only.

A simplified Ethernet encapsulation (with an optional Intercept ID) is used for all NAT traffic. When mirroring NAT traffic, the mirror-destination must be of type **ether**. The customer packet from the (outside) IP Header onwards (including the IP header) is mirrored. The operator has the configuration option of embedding the Intercept ID into the LI packet through the use of an explicit intercept-id command. Both packet formats are described below:

Figure 13 Ethernet Mirror Examples

Standard Ethernet Mirror:

Ethernet	Destination MAC Address...	
	...Destination MAC Address	Source MAC Address...
	...Source MAC Address	
H	Ethertype (IPv4 = 0x0800)	... customer packet. Ie. IPv4

Ethernet Mirror with optional Intercept ID:

Ethernet	Destination MAC Address...	
	...Destination MAC Address	Source MAC Address...
	...Source MAC Address	
LI	Ethertype (configurable)	Intercept ID...
	...Intercept ID	Ethertype (IPv4 = 0x0800)
H	... customer packet. Ie. IPv4	

OSSG539

The contents of the highlighted fields is configurable using the following CLI:

```
li
    li-source service-id
        nat
            classic-lsn-sub router name ip address
                intercept-id id
            dslite-lsn-sub router name b4 ipv6-address
                intercept-id id
            l2-aware-sub sub-ident
                intercept-id id
```

The default Ethernet-header is to use etype 0x600 and system MAC address for both source and destination address. The configurable Ethertype and Intercept ID is only added when an intercept-id is present for the subscriber in the NAT config.

2.9 Configuration Process Overview

Figure 14 shows the process to provision basic mirroring parameters.

Figure 14 Mirror Configuration and Implementation Flow





Figure 15 shows the process to provision LI parameters.

Figure 15 Lawful Intercept Configuration and Implementation Flow



OAM_15

2.10 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.
- The destination mirroring service IDs and service parameters are persistent between router (re)boots and are included in the configuration saves.

Mirror and lawful intercept source criteria configuration (defined in **debug>mirror>mirror-source** and **config>li>li-source**) is not preserved in a configuration save (admin save). Debug mirror source configuration can be saved using **admin>debug-save**. Lawful intercept source configuration can be saved using **config>li>save**.

- Subscriber based lawful intercept source criteria is persistent across creation/existence of the subscriber. Filter or sap based lawful intercept (LI) source criteria is removed from the LI source configuration if the filter entry or SAP is deleted. Applies to the 7450 ESS and 7950 SR.
- 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 or remote source. The associated mirror source is put into an operationally down mode. Mirrored packets are not transmitted out the SAP or SDP. Each mirrored packet is silently discarded. If the mirror destination is a SAP, the SAP's discard counters are incremented.
- 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, or SDP 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.

The following are lawful intercept configuration caveats.

Network management — Operators without LI permission cannot view or manage the LI data on the node nor can they view or manage the data on the Network Management platform.

LI mirroring does not allow the configuration of ports and ingress labels as a source parameter.

2.11 Configuring Service Mirroring with CLI

This section provides information about service mirroring.

Topics in this section include:

- [Mirror Configuration Overview](#)
- [Lawful Intercept Configuration Overview](#)
- [Basic Mirroring Configuration](#)
 - [Mirror Classification Rules](#)
- [Common Configuration Tasks](#)
 - [Configuring a Local Mirror Service](#)
 - [Configuring a Remote Mirror Service](#)
 - [Configuring SDPs for Mirrors and LI](#)
 - [Configuring Lawful Intercept Parameters](#)
 - [Pseudowire Redundancy for Mirror Services Configuration Example](#)
- [Service Management Tasks](#)
 - [Modifying a Local Mirrored Service](#)
 - [Deleting a Local Mirrored Service](#)
 - [Modifying a Remote Mirrored Service](#)
 - [Deleting a Remote Mirrored Service](#)

2.11.1 Mirror Configuration Overview

SR OS mirroring can be organized in the following logical entities:

- The mirror source is defined as the location where ingress or egress traffic specific to a port, SAP, MAC or IP filter, ingress label or a subscriber is to be mirrored (copied). The original frames are not altered or affected in any way.
- An SDP is used to define the mirror destination on the source router to point to a remote destination (another router).
- A SAP is defined in local and remote mirror services as the mirror destination to where the mirrored packets are sent.
- The subscriber contains hosts which are added to a mirroring service (applies to the 7450 SR and 7750 SR only).

2.11.1.1 Defining Mirrored Traffic

In some scenarios, like using VPN services or when multiple services are configured on the same port, specifying the port does not provide sufficient resolution to separate traffic. In Nokia'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/range (for example, UDP or TCP port)
- Destination port value/range (for example, UDP or TCP port)
- DiffServ Code Point (DSCP) value
- ICMP code
- ICMP type
- IP fragments
- IP option value/mask
- Single or multiple IP option fields present
- IP option fields present
- TCP ACK set/reset
- TCP SYN set/reset
- SAP ingress/egress labels

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
- Ethernet 802.2 LLC DSAP value/mask
- Ethernet 802.2 LLC SSAP value/mask
- IEEE 802.3 LLC SNAP Ethernet Frame OUI zero/non-zero value
- IEEE 802.3 LLC SNAP Ethernet Frame PID value
- SAP ingress/egress labels

2.11.2 Lawful Intercept Configuration Overview

Lawful Intercept allows the user to access and execute commands at various command levels based on profiles assigned to the user by the administrator. LI must be configured in the **config>system>security>user>access** and **config>system>security>profile** contexts. The options include FTP, SNMP, console, and LI access.

LI parameters configured in the BOF context (**li-local-save** and **li-separate**) include the ability to access LI separately than the normal administrator. As with all BOF entities, changing the BOF file during normal system operation only results in the parameter being set for the next reboot. These BOF commands are initialized to the default values, **no li-separate** and **no-li-local-save**. A system boot is necessary for any change to the **li-separate** and **li-local-save** to become effective.

Changes to the **li-separate** and **li-local-save** configurations should be made in both primary and backup CM BOF files.

At regular intervals, a LI status event is generated by the system to indicate the mode of the LI administration, time of the last reboot, and whether local save is enabled.

2.11.2.1 Saving LI Data

Depending on location and law enforcement preferences, the node can be configured to save all LI data on local media. If the operator saves this data then when starting/restarting the system the configuration file will be processed first then the LI configuration will be restarted.

When permitted to save the data, the data is encrypted and the encryption key is unique per system and is not visible to any administrator.

To save LI data locally, the option must be configured in the **bof>li-local-save** context. Enabling this option will only be applied after a system reboot.

If an LI save is permitted, then only a local save is permitted and, by default, it will be saved to Compact Flash 3 with the filename of **li.cfg**. An explicit save command under the **config>li** context must be executed to save the LI. An LI administrator with privileges to configure LI, can execute the **li.cfg** file.

2.11.2.2 Regulating LI Access

Depending on local regulations pertaining to Lawful Intercept (LI) a node can be configured to separate normal system administration tasks from tasks of a Lawful Intercept operator.

If the separation of access is not required and any administrator can manage lawful intercept or plain mirroring, then it is not necessary to configure the **li-separate** parameter in the BOF configuration. However, to ensure logical separation, the following must occur:

- An **administrator** must create a user and configure the user as LI capable (**config>system> security>user>access** context). Furthermore, the **administrator** must assure that both CLI and SNMP access permission is granted for the LI operator.
- Finally, before turning the system into two separate administration domains, the CLI user must be granted a profile that limits the LI operator to those tasks relevant to the job (**config>system> security>profile>li** context).

It is important to remember that the LI operator is the only entity who can grant LI permission to any other user once in **li-separate** mode.

Provided the above procedure is followed, the LI administrator must decide whether to allow the LI (source) configuration to be saved onto local media. This is also subject to local regulations.

At this point, the BOF file can be configured with the **li-separate** and **li-local-save** parameters. If the local save is not configured then the LI information must be reconfigured after a system reboot.

Assuming **li-separate** is configured, the node should be rebooted to activate the **separate** mode. At this point the system administrators without LI permission cannot modify, create or view any LI- specific configurations. In order for this to occur, the BOF file must be reconfigured and the system rebooted. This, combined with other features prohibits an unauthorized operator from modifying the administrative separation without notifying the LI administrator.

The following example shows an SNMP configuration with views, access groups, and attempts parameters.

```
A:ALA-23>config>system>security>snmp# info detail
-----
view iso subtree 1
  mask ff type included
exit
view no-security subtree 1
  mask ff type included
exit
```


```
view no-security subtree 1.3.6.1.6.3
  mask ff type excluded
exit
view no-security subtree 1.3.6.1.6.3.10.2.1
  mask ff type included
exit
view no-security subtree 1.3.6.1.6.3.11.2.1
  mask ff type included
exit
view no-security subtree 1.3.6.1.6.3.15.1.1
  mask ff type included
exit
...
access group "snmp-li-ro" security-model usm security-
level <security level>
context "li" read "li-view" notify "iso"
access group "snmp-li-rw" security-model usm security-
level <security level>
context "li" read "li-view" write "li-view" notify "iso"
  attempts 20 time 5 lockout 10
...
-----
A:ALA-23>config>system>security>snmp#
```

The following example shows a user account configuration.

```
A:ALA-23>config>system>security# info
-----
...
user "liuser"
  access console snmp li
  console
    no member "default"
    member "liprofile"
  exit
  snmp
    authentication md5 <auth-key> privacy des <priv-key>
    group "snmp-li-rw"
  exit
exit
...
-----
A:ALA-23>config>system>security#
```

2.11.2.2.1 LI User Access

By default, LI user access is limited to those commands that are required to manage LI functionality. If a user is granted permission to access other configuration and operational data, then this must be explicitly configured in the user profile of the LI operator in the **config>system>security>profile>entry>match command-string** context. [Figure 16](#) shows the work flow to set an LI operator.

Figure 16 Creating an LI Operator Account

OSSG264

2.11.2.2.2 LI Source Configuration

Filter configuration is accessible to both the LI operator and regular system administrators. If the content of a filter list that is subject to an LI operation and if a filter (included in the filter list) is used by an LI operator, its contents cannot be modified unless the **li-filter-lock-state** is unlocked, see [Configurable Filter Lock for Lawful Intercept](#). If an attempt is made, then an LI event is generated. Only one mirror source, which can contain one or many li-source entries, can be attached to one mirror destination service. LI takes priority over debug mirror sources. So if a debug mirror source (for example, 10) exists and an LI mirror source is created with same ID 10, then the debug mirror source is silently discarded.

In the configuration, when an LI operator specifies that a given entry must be used as an LI entry then this fact is hidden from all non-LI operators. Modification of a filter entry is not allowed if it is used by LI, see [Configurable Filter Lock for Lawful Intercept](#). However, an event is generated, directed to the LI operator, indicating that the filter has been compromised.

Standard mirroring (non-LI) has a lower priority than LI instantiated mirroring. If a mirror source parameter (for example, SAP 1/1/1) exists and the same parameter is created in an LI source, the parameter is silently deleted from the debug mirror source.

The following order applies for both ingress and egress traffic:

- Port mirroring (debug only)
- SAP mirroring (debug or LI)
- Subscriber mirroring (debug or LI) for the 7450 ESS and 7750 SR
- Filter mirroring (debug or LI)

For frames from network ports:

- Port mirroring (debug only)
- Label mirroring (debug only, ingress only)
- Filter mirroring (debug or LI)

Filters can be created by all users that have access to the relevant CLI branches.

Once an LI mirror source using a given service ID is created and is in the **no shutdown** state, the corresponding mirror destination on the node cannot be modified (including **shutdown/no shutdown** commands) or deleted.

In the **separate** mode, the anonymity of the source is protected. Once source criterion is attached to the LI source, the following applies:

- In SAP configurations, only modifications that stop the flow of LI data while the customer receives data is blocked unless the `li-filter-lock-state` is unlocked, see [Configurable Filter Lock for Lawful Intercept](#).
- In filter configurations, if a filter entry is attached to the LI source, modification and deletion of both the filter and the filter entry are blocked.

2.11.2.3 Configurable Filter Lock for Lawful Intercept

With the default Lawful Intercept configuration, when a filter entry is used as a Lawful Intercept (LI) mirror source criteria/entry, all subsequent attempts to modify the filter are then blocked to avoid having the LI session impacted by a non-LI user.

A configurable LI parameter allows an a LI user to control the behavior of filters when they are used for LI.

Configuration of the **li-filter-lock-state** allows an operator to control whether modifications to filters that are being used for LI are allowed by no users, all users or li users only.

2.11.2.4 LI MAC Filter Configuration

Although normal MAC filter entries (configured under **config>filter>mac-filter**) can be referenced in an **li-source**, there is also the option to configure and use special-purpose Lawful Intercept MAC filters.

LI MAC filters are configured in the protected **config>li** CLI branch.

LI MAC filters are associated by configuration with normal MAC filters, and entries created in the LI MAC filters are inserted into the associated normal MAC filter before the filter is downloaded to the data plane hardware and applied. The combined filter list is not visible to any users which maintains a separation between LI operators and operators doing other normal filter configuration work (e.g. interface ACLs).

A configurable **li-filter-block-reservation** is used to reserve a range of entries in the normal filter into which the LI entries are inserted.

2.11.2.5 LI Logging

A logging collector is supported in addition to existing main, security, change, and debug log collectors. LI log features include the following:

- Only visible to LI operators (such as show command output).
- Encrypted when transmitted (SNMPv3).
- Logging ability can only be created, modified, or deleted by an LI operator.
- The LI user profile must include the ability to manage the LI functions.

2.12 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 or SDP.

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, ingress label, IP filter or MAC filter) specified.

The following example shows a configuration of a local mirrored service where the source and destinations are on the same device (ALA-A).

```
*A:ALA-A>config>mirror# info
-----
      mirror-dest 103 create
          sap 2/1/25:0 create
egress
          qos 1
          exit
      exit
      no shutdown
      exit
-----
*A:ALA-A>config>mirror#
```

The following examples shows a mirror source configuration:

```
*A:ALA-A>debug>mirror-source# show debug mirror
debug
      mirror-source 103
          port 2/1/24 egress ingress
no shutdown
      exit
exit
*A:ALA-A>debug>mirror-source# exit
```

The following example shows a configuration of a remote mirrored service where the source is a port on ALA-A and the destination is a SAP is on ALA-B:

```
*A:ALA-A>config>mirror# info
-----
      mirror-dest 1000 create
      spoke-sdp 2:1 egr-svc-label 7000
      no shutdown
      exit
-----
*A:ALA-A>config>mirror# exit all
*A:ALA-A# show debug
debug
      mirror-source 1000
      port 2/1/2 egress ingress
no shutdown
      exit
exit
*A:ALA-A#

*A:ALA-B>config>mirror# info
-----
      mirror-dest 1000 create
      remote-source
      far-end 10.10.10.104 ing-svc-label 7000
      exit
      sap 3/1/2:0 create
egress
      qos 1
      exit
      exit
      no shutdown
      exit
-----
*A:ALA-B>config>mirror#
```

2.12.1 Mirror Classification Rules

Nokia'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](#)
- [Ingress label](#)
- [Subscriber](#)

2.12.1.1 Port

The port command associates a port to a mirror source. The port is identified by the port ID.

The following shows the *port-id* syntax for the port command:

<i>port-id:</i>	slot/mda/port[.channel]	
eth-sat-id	esat-id/slot/port	
	esat	keyword
	<i>id</i>	1 to 20
pxc-id	pxc-id.sub-port	
	pxc	keyword
	<i>id</i>	1 to 64
	<i>sub-port</i>	a, b
aps-id	aps-group-id[.channel]	
	aps	keyword
	<i>group-id</i>	1 to 64
bundle-type-slot/mda.bundle-num		
	bundle	keyword
	<i>type</i>	ima
	<i>bundle-num</i>	1 to 128
ccag-id	- ccag-id.path-id[cc-type]:cc-id	
	ccag	keyword
	<i>id</i>	1 o 8
	<i>path-id</i>	a,b
	<i>cc-type</i>	.sap-net, .net-sap
	<i>cc-id</i>	0 to 4094
lag-id	1 to 800	
egress	keyword	
ingress	keyword	



Note: On the 7950 XRS, the XMA ID takes the place of the MDA.

The defined port can be an Ethernet or Frame Relay port, a SONET/SDH path, a multilink bundle, a TDM channel, a Cross Connect Aggregation Group (CCAG), or a Link Aggregation Group (LAG) ID. If the port is a SONET/SDH or TDM channel, the channel ID must be specified to identify which channel is being mirrored. When a LAG ID is given as the port ID, mirroring is enabled on all ports making up the LAG. Ports that are circuit-emulation (CEM) and PPP bundle groups cannot be used in a mirror source (applies to the 7750 SR). Note, Frame Relay and SONET/SDH apply to the 7450 ESS and 7750 SR, and multilink bundle and TDM channel apply to the 7750 SR.

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

Table 3 Mirror Source Port Requirements

Port Type	Port Mode	Port Encap Type
faste/gige/xgige ethernet	access	dot1q, null, qinq
faste/gige/xgige ethernet	network	dot1q, null
SONET (clear/deep channel)	access	bcp-null, bcp-dot1q, ipcp
TDM (clear/deep channel)	access	bcp-null, bcp-dot1q, ipcp

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

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

2.12.1.2 SAP

More than one SAP can be associated within a single mirror-source. Each SAP has its own ingress and egress parameter keywords 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 { [egress] [ingress] }`

Example: `*A:ALA-A>debug>mirror-source# sap 2/1/4:100 ingress
egress`

`or debug>mirror-source# port 2/2/1.sts12 ingress`

2.12.1.3 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

2.12.1.4 IP filter

IP filters are configured in the **config>filter>ip-filter** or **config>filter>ipv6-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. Egress mirrored packets are mirrored to the mirror destination after all egress packet modifications.

CLI Syntax: debug>mirror-source# ip-filter *ip-filter-id* entry *entry-id*
 [*entry-id* ...]
 debug>mirror-source# ipv6-filter *ipv6-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.

2.12.1.5 Ingress label

The **ingress-label** command is used to mirror ingressing MPLS frames with the specified MPLS labels. The ingress label must be at the top of the label stack and can only be mirrored to a single mirror destination. If the same label is defined with multiple mirror destinations, an error is generated and the original mirror destination does not change. The **ingress-label** allows packets matching the ingress label to be duplicated (mirrored) and forwarded to the mirror destination. The ingress label has to be active before it can be used as mirror source criteria. If the ingress label is not used in the router, the mirror source will remove the ingress label automatically.

CLI Syntax: debug>mirror-source# ingress-label *label* [*label...*]

Example: *A:ALA-A>debug>mirror-source# ingress-label 103000
 1048575

2.12.1.6 Subscriber

The subscriber command is used to add hosts of a subscriber to a mirroring service. This command applies to the 7450 ESS and 7750 SR only.

CLI Syntax: debug>mirror-source# subscriber *sub-ident-string*
 [*sap...*]


2.13 Common Configuration Tasks

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

Each local mirrored service (Figure 17) (within the same router) requires the following configurations:


- Step 1.** Specify mirror destination (SAP).
- Step 2.** Specify mirror source (port, SAP, IP filter, MAC filter, ingress label, subscriber).

Figure 17 Local Mirrored Service Tasks



Each remote mirrored service (Figure 18) (across the network core) requires the following configurations:

- Step 1.** Define the remote destination (SDP)
- Step 2.** Identify the remote source (the device allowed to mirror traffic to this device)
- Step 3.** Specify the mirror destination (SAP)
- Step 4.** Specify mirror source (port, SAP, IP filter, MAC filter, ingress label, subscriber)

Figure 18 Remote Mirrored Service Configuration Example

2.13.1 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, in the same mirror-source an entire port X could be mirrored at the same time as packets matching a filter entry applied to SAP Y could be mirrored. A filter must be applied to the SAP or interface if only specific packets are to be mirrored. Note that slice-size is not supported by CEM encap-types or IP-mirroring (applies to the 7750 SR and 7950 XRS).

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, the forwarding class, and the size of the packet.

The following output shows an example of a local mirrored service. On ALA-A, mirror service 103 is mirroring traffic matching IP filter 2, entry 1 as well as egress and ingress traffic on port 2/1/24 and sending the mirrored packets to SAP 2/1/25:

```
*A:ALA-A>config>mirror# info
-----
mirror-dest 103 create
  sap 2/1/25:0 create
egress
  qos 1
  exit
exit
```

```

        no shutdown
        exit
-----
*A:ALA-A>config>mirror#

```

The following output shows debug mirroring information:

```

*A:ALA-A>debug>mirror-source# show debug mirror
debug
    mirror-source 103
no shutdown
    port 2/1/24 egress ingress
ip-filter 2 entry 1
    exit
exit
*A:ALA-A>debug>mirror-source# exit

```

Note that the ip-filter and entry referenced by the mirror-source must exist and must be applied to an object in order for traffic to be mirrored:

```

*A:ALA-A>config>service>vprn>if# info

```

```

-----
        sap 1/1/3:63 create
        ingress
            filter ip 2
        exit
    exit
-----

```

2.13.2 Configuring SDPs for Mirrors and LI

This section provides a brief overview of the tasks that must be performed to configure SDPs and provides the CLI commands. For more information about service configuration, refer to the Services Guide.

Consider the following SDP characteristics:

- Configure GRE, MPLS, MPLS-TP or L2TPv3 SDPs.
- Each distributed service must have an SDP defined for every remote SR to provide Epipe, VPLS, or mirrored services.
- A distributed service must be bound to an SDP. By default, no SDP is associated with a service. Once an SDP is created, services can be associated to that SDP.
- An SDP is not specific to any one service or any type of service. An SDP can have more than one service bound to it.

- When using L2TPv3, MPLS-TP or LDP IPv6 LSP SDPs in a remote mirroring solution, configure the destination node with **remote-src>spoke-sdp** entries. For all other types of SDPs use **remote-src>far-end** entries.
- In order to configure an MPLS SDP, LSPs must be configured first and then the LSP-to-SDP association must be explicitly created.

To configure a basic SDP, perform the following steps:

- Step 1.** Select an originating node.
- Step 2.** Create an SDP ID.
- Step 3.** Select an encapsulation type.
- Step 4.** Select the far-end node.

To configure the return path SDP, perform the same steps on the far-end router.

- Step 1.** Select an originating node.
- Step 2.** Create an SDP ID.
- Step 3.** Select an encapsulation type.
- Step 4.** Select the far-end node.

Use the following CLI syntax to create an SDP and select an encapsulation type. If you do not specify GRE or MPLS, the default encapsulation type is GRE.



Note: When you specify the far-end IP address, you are creating the tunnel. In essence, you are creating the path from Point A to Point B. Use the `show service sdp` command to display the qualifying SDPs.

CLI Syntax:

```
config>service# sdp sdp-id [gre | mpls] create
description description-string
far-end ip-addr
lsp lsp-name [lsp-name]
path-mtu octets
no shutdown
keep-alive
    hello-time seconds
    hold-down-time seconds
    max-drop-count count
    message-length octets
    no shutdown
```

On the mirror-source router, configure an SDP pointing toward the mirror-destination router (or use an existing SDP).

On the mirror-destination router, configure an SDP pointing toward the mirror-source router (or use an existing SDP).

The following example shows SDP configurations on both the mirror-source and mirror-destination routers.

```
*A:ALA-A>config>service# info
-----
sdp 1 create
      description "to-10.10.10.104"
      far-end 10.10.10.104
      no shutdown
      exit
-----
*A:ALA-A>config>service#


*A:ALA-B>config>service# info
-----
      sdp 4 create
      description "to-10.10.10.103"
      far-end 10.10.10.103
      no shutdown
      exit
-----
*A:ALA-B>config>service#
```

2.13.3 Configuring a Remote Mirror Service

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

When using L2TPv3, MPLS-TP or LDP IPv6 LSP spoke SDPs in a remote mirroring solution, configure the destination node with **remote-src>spoke-sdp** entries. For all other types of SDPs use **remote-src>far-end** entries.

[Figure 19](#) shows 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 4/1/58 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. The slice size limits the size of the stream of packet through the router and the core network.

Figure 19 Remote Mirrored Service Tasks

OSSG030

The following example shows 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.

```
*A:ALA-A>config>mirror# info
-----
mirror-dest 1216 create
  description "Receiving mirror traffic from .91"
  remote-source
    far-end 10.10.0.91 ing-svc-label 5678
  exit
  sap 1/1/58:0 create
    egress
      qos 1
    exit
  exit
  no shutdown
exit
-----
*A:ALA-A>config>mirror#
```

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

```
*A:ALA-B>config>mirror># info
-----
mirror-dest 1216 create
  description "Sending mirrored traffic to .92"
  fc h1
  spoke-sdp 4:60 create
    egress
  vc-label 5678
  exit
  no shutdown
  exit
  slice-size 128
  no shutdown
  exit
-----
```

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

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

```
*A:ALA-B# show debug mirror
debug
  mirror-source 1216
    port 1/1/60 egress ingress
no shutdown
exit
exit
*A:ALA-B#
```

The following example shows 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 "GRE-10.10.0.91"
far-end 10.10.0.01
no shutdown
-----
*A:ALA-A>config>service>sdp#

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

2.13.4 Configuring an ATM Mirror Service

The ATM Mirror Service applies to the 7750 SR only.

Configure a local ATM mirror service at PE1:

Example:

```
config>mirror# mirror-dest 1 type atm-sdu create
config>mirror>mirror-dest# sap 1/2/1:1/101 create
config>mirror>mirror-dest>sap# no shutdown
config>mirror>mirror-dest>sap# exit all
# debug
debug# mirror-source 1
debug>mirror-source# sap 2/1/1/:0/100 ingress
```

Configure a remote ATM mirror service at PE1:

Example:

```
config>mirror# mirror-dest 1 type atm-sdu create
```

```

config>mirror>mirror-dest# spoke-sdp 1:20
config>mirror>mirror-dest# exit all
# debug

debug# mirror-source 1
debug>mirror-source# sap 2/1/1/:0/100 ingress

```

Configure a remote ATM mirror service at PE2:

Example:

```

config>mirror# mirror-dest 1 type atm-sdu create
config>mirror>mirror-dest# remote-source
config>mirror>mirror-dest>remote-source# far-end
10.10.10.10
config>mirror>mirror-dest>remote-source# exit
config>mirror>mirror-dest# sap 1/2/1:1/101 create

```

2.13.5 Configuring Lawful Intercept Parameters

The following example shows an LI source configuration and LI log configuration examples:

```

A:ALA-48>config# info
#-----
...
(LI Source Config)
  li-source 1
    sap 1/5/5:1001 egress ingress
    no shutdown
  exit
  li-source 2
    subscriber "test" sla-profile "test" fc l2 ingress egress
    no shutdown
  exit
  li-source 3
    mac-filter 10 entry 1
    no shutdown
  exit
  li-source 4
    ip-filter 11 entry 1
    no shutdown
  exit
...
(LI Log Config)
  log-id 1
    filter 1
    from li
    to session
  exit
  log-id 11
    from li
    to memory

```



```


exit
log-id 12
  from li
  to snmp
exit
...
#-----
A:ALA-48>config#

```

2.13.6 Pseudowire Redundancy for Mirror Services Configuration Example

A configuration based on [Figure 20](#) is described in this section.

Figure 20 State Engine for Redundant Service to a Redundant Mirror Service



The mirror traffic needs to be forwarded from configured debug mirror-source together with mirror-dest/remote-source (ICB or non-ICB) to either SAP endpoint or SDP endpoint.

A SAP endpoint is an endpoint with a SAP and with or without an additional ICB spoke. An SDP endpoint is an endpoint with regular and ICB spokes.

Only one tx-active will be chosen for either SAP endpoint or SDP endpoint. Traffic ingressing into a remote-source ICB will have only ingressing traffic while an ICB spoke will have only egressing traffic.

The ingressing traffic to a remote-source ICB cannot be forwarded out of another ICB spoke.

The following example shows a high level summary of a configuration; it is not intended to be syntactically correct:

```
Node A:
config mirror mirror-dest 100
endpoint X
sdp to-C endpoint X
sdp to-D endpoint X
sdp to-B endpoint X icb // connects to B's remote-source IP-A, traffic A->B only
remote-source IP-B icb // connects to B's sdp to-A, traffic B->A only

Node B:
config mirror mirror-dest 100
endpoint X
sdp to-C endpoint X
sdp to-D endpoint X
sdp to-A endpoint X icb // connects to A's remote-source IP-B, traffic B->A only
remote-source IP-A icb // connects to Node A's sdp to-B, traffic A->B only

Node C:
config mirror mirror-dest 100
endpoint X
sap lag-1:0 endpoint X
sdp to-D endpoint X icb // connects to D's remote-source IP-C, traffic C->D only
remote-source IP-A
remote-source IP-B
remote-source IP-D icb // connects to D's sdp to-C, traffic D->C only

Node D:
config mirror mirror-dest 100
endpoint X
sap lag-1:0 endpoint X
sdp to-C endpoint X icb // connects to C's remote-source IP-D, traffic D->C only
remote-source IP-A
remote-source IP-B
remote-source IP-C icb // connects to C's sdp to-D, traffic C->D only
```

2.14 Service Management Tasks

This section describes the following service management tasks:

- [Modifying a Local Mirrored Service](#)
- [Deleting a Local Mirrored Service](#)
- [Modifying a Remote Mirrored Service](#)
- [Deleting a Remote Mirrored Service](#)

2.14.1 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 shows the 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 3/1/5:0 create
                  config>mirror>mirror-dest>sap$ exit
                  config>mirror>mirror-dest# fc be
                  config>mirror>mirror-dest# slice-size 128
                  config>mirror>mirror-dest# no shutdown

                  debug# mirror-dest 103
                  debug>mirror-source# no port 2/1/24 ingress egress
                  debug>mirror-source# port 3/1/7 ingress egress
```

The following output shows the local mirrored service modifications:

```
*A:ALA-A>config>mirror# info
-----
mirror-dest 103 create
          no shutdown
          fc be
          remote-source
          exit
          sap 3/1/5:0 create
egress
          qos 1
          exit
          exit
          slice-size 128
          exit

*A:ALA-A>debug>mirror-source# show debug mirror
debug
      mirror-source 103
          no shutdown
          port 3/1/7 egress ingress
exit
*A:ALA-A>debug>mirror-source#
```

2.14.2 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 shows the 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
```

2.14.3 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 shows the 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# spoke-sdp 4:60 egress vc-
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 10.10.10.3 ing-svc-label 3500
    exit
    sap 2/1/15:0 create
egress
    qos 1
    exit
    exit
    no shutdown
exit

A:SR3>config>mirror# info
-----
mirror-dest 104 create
    spoke-sdp 4:60 egress vc-label 3500
    no shutdown
    exit
-----

A:SR3>config>mirror#

A:SR3# show debug mirror
debug
    mirror-source 104
        no shutdown
        port 5/1/2 egress ingress
exit
    exit
A:SR3#

```

2.14.4 Deleting a Remote Mirrored Service

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

Mirror destinations must be shut down first before they can be 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
```

In the example, 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#
```

2.15 Mirror Service Configuration Command Reference

2.15.1 Command Hierarchies

- [Mirror Configuration Commands](#)
- [Lawful Intercept Commands](#)

2.15.1.1 Mirror Configuration Commands

```

config
— mirror
— mirror-dest service-id [type mirror-type] [create]
— no mirror-dest service-id
— description description-string
— no description
— [no] enable-port-id
— encap
— layer-3-encap {ip-udp-shim | ip-gre} [create]
— no layer-3-encap
— direction-bit
— no direction-bit
— router {router-instance | service-name service-name}
— no router
— gateway [create]
— no gateway
— ip src ip-address dest ip-address
— no ip
— udp src udp-port dest udp-port
— no udp
— endpoint endpoint-name [create]
— no endpoint endpoint-name
— description description-string
— no description
— revert-time {revert-time | infinite}
— no revert-time
— fc fc-name
— no fc
— isa-aa-group aa-group-id traffic-direction
— [no] remote-source
— far-end ip-address [vc-id vc-id][ing-svc-label ing-vc-label | tldp] [icb]
— no far-end ip-address
— spoke-sdp sdp-id:vc-id [create] [no-endpoint]
— spoke-sdp sdp-id:vc-id [create] endpoint <name> [icb]
— no spoke-sdp sdp-id:vc-id

```

```

— [no] control-channel-status
— [no] acknowledgment
— refresh-timer value
— no refresh-timer
— request-timer timer1 retry-timer timer2 [timeout-
    multiplier multiplier]
— no request-timer
— [no] shutdown
— [no] control-word
— egress
— vc-label egress-vc-label
— no vc-label [egress-vc-label]
— ingress
— l2tpv3
— cookie <cookie1> <cookie2>
— vc-label ingress-vc-label
— no vc-label [ingress-vc-label]
— [no] pw-path-id
— agi route-identifier
— no agi
— saii-type2 global-id:node-id:ac-id
— no saii-type2
— taii-type2 global-id:node-id:ac-id
— no taii-type2
— [no] shutdown
— sap sap-id [create] [no-endpoint]
— sap sap-id [create] endpoint name
— no sap
— cem
— packet jitter-buffer milliseconds [payload-size bytes]
— packet payload-size bytes
— no packet bytes
— [no] rtp-header
— egress
— ip-mirror
— sa-mac ieee-address da-mac ieee-address
— no sa-mac
— qos policy-id
— no qos
— service-name service-name
— no service-name
— [no] shutdown
— slice-size bytes
— no slice-size
— spoke-sdp sdp-id:vc-id [create] [no-endpoint]
— spoke-sdp sdp-id:vc-id [create] endpoint name [icb]
— no spoke-sdp sdp-id:vc-id
— [no] control-channel-status
— [no] acknowledgment
— refresh-timer value
— no refresh-timer
— request-timer timer1 retry-timer timer2 [timeout-multiplier
    multiplier]
— no request-timer

```


- [no] **shutdown**
- [no] **control-word**
- **egress**
 - **vc-label** *egress-vc-label*
 - **no vc-label** [*egress-vc-label*]
 - **l2tpv3**
 - **cookie** <*cookie*>
- **ingress**
 - **vc-label** *ingress-vc-label*
 - **no vc-label** [*ingress-vc-label*]
- **precedence** *precedence-value* | **primary**
- **no precedence**
- [no] **pw-path-id**
 - **agi** *route-identifier*
 - **no agi**
 - **saii-type2** *global-id:node-id:ac-id*
 - **no saii-type2**
 - **taii-type2** *global-id:node-id:ac-id*
 - **no taii-type2**
- [no] **shutdown**
- [no] **shutdown**

2.15.1.2 Lawful Intercept Commands

```
config
  — li
    — li-filter
      — li-mac-filter filter-name [create]
      — no li-mac-filter
        — description description-string
        — no description
        — entry li-entry-id
        — no entry
          — description description-string
          — no description
          — match [frame-type {802dot3 | 802dot2-llc | 802dot2-snap ethernet_II}]
          — no match
            — dst-mac ieee-address [mask]
            — no dst-mac
            — src-mac ieee-address [mask]
            — no src-mac
      — li-ip-filter filter-name [create]
      — no li-ip-filter
        — description description-string
        — no description
        — entry li-entry-id
        — no entry
          — description description-string
          — no description
          — match [protocols protocol-id]
```

```

      — no match
        — dst-ip ipv4-address/mask
        — no dst-ip
        — src-ip ipv4-address/mask
        — no src-ip
— li-ipv6-filter filter-name [create]
— no li-ipv6-filter
  — description description-string
  — no description
  — entry li-entry-id
  — no entry
    — description description-string
    — no description
    — match [next-header next-header]
    — no match
      — dst-ip ipv6-address/prefix-length
      — no dst-ip
      — src-ip ipv6-address/prefix-length
      — no src-ip
— li-filter-associations
  — li-mac-filter li-mac-filter-name
    — mac-filter mac-filter-id
  — li-ip-filter li-ip-filter-name
    — ip-filter ip-filter-id
  — li-ipv6-filter li-ipv6-filter-name
    — ipv6-filter ipv6-filter-id
— li-filter-block-reservation
  — [no] li-reserved-block block-name [create]
    — description description-string
    — [no] start-entry entry-id count count
    — [no] ip-filter filter-id
    — [no] ipv6-filter filter-id
    — [no] mac-filter filter-id
— [no] li-filter-lock-state {locked | unlocked-for-li-users | unlocked-for-all-users}
— li-source service-id
  — ip-filter ip-filter-id [entry entry-id...] [intercept-id id] [session-id id]
  — no ip-filter ip-filter-id
  — ipv6-filter ipv6-filter-id [entry entry-id...] [intercept-id intercept-id...] [session-id session-id...]
  — no ipv6-filter ipv6-filter-id
  — li-ip-filter ip-filter-id [entry entry-id...] [intercept-id id] [session-id id]
  — no li-ip-filter ip-filter-id
  — li-ipv6-filter ipv6-filter-id [entry entry-id...] [intercept-id intercept-id...] [session-id session-id...]
  — no li-ipv6-filter ipv6-filter-id
  — mac-filter mac-filter-id entry entry-id [entry-id...] [intercept-id id] [session-id id]
  — no mac-filter mac-filter-id
— nat
  — [no] classic-lsn-sub router router-instance ip ip-address
    — intercept-id id
    — no intercept-id
    — session-id id
    — no session-id

```

- [no] **dslite-lsn-sub** router *router-instance* **b4** *ipv6-prefix*
 - **intercept-id** *id*
 - **no intercept-id**
 - **session-id** *id*
 - **no session-id**
- **ethernet-header** [**da** *ieee-address*] [**sa** *ieee-address*] [**etype** *ethertype*]
- **no ethernet-header**
- [no] **l2-aware-sub** *sub-ident-string*
 - **intercept-id** *id*
 - **no intercept-id**
 - **session-id** *id*
 - **no session-id**
- [no] **nat64-lsn-sub** router *router-instance* **ip** *ipv6-prefix*
 - **intercept-id** *id*
 - **no intercept-id**
 - **session-id** *id*
 - **no session-id**
- **sap** *sap-id* {[**ingress**] [**egress**]} [**intercept-id** *id*] [**session-id** *id*]
- **no sap** *sap-id*
- [no] **shutdown**
- **subscriber** *sub-ident-string* [**sap** *sap-id*] [**ip** *ip-address*] [**mac** *ieee-address*]
 - [**sla-profile** *sla-profile-name*] [**fc** {[**be**] [**l2**] [**af**] [**l1**] [**h2**] [**ef**] [**h1**] [**nc**]}]
 - {[**ingress**] [**egress**]} [**intercept-id** *id*] [**session-id** *id*]
- **no subscriber** *sub-ident-string*
- **wlan-gw**
 - [no] **dsm-subscriber** **mac** *xx:xx:xx:xx:xx:xx* or *xx-xx-xx-xx-xx-xx*
 - **intercept-id** [1..4294967295]
 - **no intercept-id**
 - **session-id** [1..4294967295]
 - **no session-id**
- **log**
 - [no] **log-id** *log-id*
 - **description** *description-string*
 - **no description**
 - **filter** *filter-id*
 - **no filter**
 - **from** {[**li**]}
 - **no from**
 - [no] **shutdown**
 - **time-format** {**local** | **utc**}
 - **to memory** [*size*]
 - **to session**
 - **to snmp** [*size*]
- **mobility-gateway**
 - **df-peer** *df-peer-id* **df2-addr** *ip-address* **df2-port** *port* **df3-addr** *ip-address* **df3-port** *port*
 - **no df-peer** *df-peer-id*
 - **local-interface** *ip-address* [**router** *router-instance*]
 - **no local-interface**
 - **operator** *op-id*
 - **no operator**
 - **target** *target-type* **id** *string* **intercept** *intercept* **peer** *df-peer-id* [**liid** *li-identifier*]
 - **no target** *target-type* **id** *string*
 - [no] **x2-iri-qos** **dscp** {*dscp-value* | *dscp-name*}

- [no] **x3-cc-qos** **dscp** {*dscp-value* | *dscp-name*}
- **x3-transport** {*tcp* | *udp*} **ulic-header** {*v0* | *v1*}
- **save**

The following commands are also described in the *Basic System Configuration Guide*.

- ```
config
 — bof
 — [no] li-local-save
 — [no] li-separate
```

The following commands are also described in the *7450 ESS, 7750 SR, and 7950 XRS System Management Guide*.

- ```
config
  — system
    — security
      — user
        — [no] access [ftp] [snmp] [console] [li]
        — [no] profile user-profile-name
        — [no] li
```

2.15.2 Command Descriptions

2.15.2.1 Generic Commands

description

Syntax	description <i>description-string</i> no description
Context	config>mirror>mirror-dest config>li>log>log-id config>li>li-filter>li-mac-filter config>li>li-filter>li-mac-filter>entry config>li>li-filter>li-ip-filter config>li>li-filter>li-ip-filter>entry config>li>li-filter>li-ipv6-filter config>li>li-filter>li-ipv6-filter>entry config>li>li-filter-block-reservation>li-reserved-block
Description	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 from the configuration.

Default	There is no default description associated with the configuration context.
Parameters	<i>description-string</i> — 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 debug>mirror-source config>mirror>mirror-dest>spoke-sdp>egress config>li>li-source config>li>log>log-id
Description	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. 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 or remote source router. The associated mirror source is put into an operationally down mode. Mirrored packets are not transmitted out of the SAP or SDP. Each mirrored packet is silently discarded. If the mirror destination is a SAP, the SAP's discard counters are incremented. 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 or SDP 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. If the **remote-source** command has been executed on the **mirror-dest** associated with the shutdown **mirror-source**, mirroring continues for remote sources.

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.

2.15.2.2 Mirror Destination Configuration Commands

enable-port-id

Syntax	[no] enable-port-id
Context	config>mirror>mirror-dest
Description	This command includes the mirrored packet system's port-id. The system port ID can be used to identify which port the packet was received or sent on. Inclusion of the port-id is only supported for mirror-dest type ppp.
Default	no enable-port-id

endpoint

Syntax	endpoint <i>endpoint-name</i> [create] no endpoint <i>endpoint-name</i>
Context	config>mirror>mirror-dest config>mirror>mirror-dest>sap config>mirror>mirror-dest>sdp
Description	A mirror service supports two implicit endpoints managed internally by the system. The following applies to endpoint configurations. Up to two (2) named endpoints may be created per service mirror/LI service. The endpoint name is locally significant to the service mirror/LI service. • Objects (SAPs or SDPs) may be created on the service mirror/LI with the following limitations: – two implicit endpoint objects (without explicit endpoints defined) – one implicit and multiple explicit object with the same endpoint name – multiple explicit objects each with one of two explicit endpoint names • All objects become associated implicitly or indirectly with the implicit endpoints 'x' and 'y'. • Objects may be created without an explicit endpoint defined.

- Objects may be created with an explicit endpoint defined.
- Objects without an explicit endpoint may have an explicit endpoint defined without deleting the object.
- Objects with an explicit endpoint defined may be dynamically moved to another explicit endpoint or may have the explicit endpoint removed.

Creating an object without an explicit endpoint:

- If an object on a mirror/LI service has no explicit endpoint name associated, the system attempts to associate the object with implicit endpoint 'x' or 'y'.
- The implicit endpoint cannot have an existing object association.
- If both 'x' and 'y' are available, 'x' will be selected.
- If an 'x' or 'y' association cannot be created, the object cannot be created.

Creating an object with an explicit endpoint name:

- The endpoint name must exist on the mirror/LI service.
- If this is the first object associated with the endpoint name:
 - the object is associated with either implicit endpoint 'x' or 'y'
 - the implicit endpoint cannot have an existing object associated
 - if both 'x' and 'y' are available, 'x' will be selected
 - if 'x' or 'y' is not available, the object cannot be created
 - the implicit endpoint is now associated with the named endpoint
- if this is not the first object associated with the endpoint name:
 - the object is associated with the named endpoint's implicit association

Changing an objects implicit endpoint to an explicit endpoint name

- If the explicit endpoint name is associated with an implicit endpoint, the object is moved to that implicit endpoint
- If the object is the first to be associated with the explicit endpoint name:
 - the object is associated with either implicit endpoint 'x' or 'y'
 - the implicit endpoint cannot have an existing object associated (except this one)
 - if both 'x' and 'y' are available, 'x' will be selected
 - if 'x' or 'y' is not available, the object cannot be moved to the explicit endpoint
 - if moved, the implicit endpoint is now associated with the named endpoint

Changing an objects explicit endpoint to another explicit endpoint name

- If the new explicit endpoint name is associated with an implicit endpoint, the object is moved to that implicit endpoint
- If the object is the first to be associated with the new explicit endpoint name:
 - the object is associated with either implicit endpoint 'x' or 'y'
 - the implicit endpoint cannot have an existing object associated (except this one)
 - if both 'x' and 'y' are available, 'x' will be selected

- if 'x' or 'y' is not available, the object cannot be moved to the new endpoint
- if moved, the implicit endpoint is now associated with the named endpoint

An explicitly named endpoint can have a maximum of one SAP and one ICB. Once a SAP is added to the endpoint, only one more object of type ICB sdp is allowed. The ICB sdp cannot be added to the endpoint if the SAP is not part of a MC-LAG instance. Conversely, a SAP which is not part of a MC-LAG instance cannot be added to an endpoint which already has an ICB sdp.

An explicitly named endpoint which does not have a SAP object can have a maximum of four SDPs which can include any of the following: a single primary SDP, one or many secondary SDPs with precedence, and a single ICB SDP.

The user can only add a SAP configured on a MC-LAG instance to this endpoint. Conversely, the user will not be able to change the mirror service type away from mirror service without first deleting the MC-LAG SAP.

The **no** form of the command removes the association of a SAP or a sdp with an explicit endpoint name. Removing an objects explicit endpoint association:

- The system attempts to associate the object with implicit endpoint 'x' or 'y'.
- The implicit endpoint cannot have an existing object association (except this one).
- If both 'x' and 'y' are available, 'x' will be selected.
- If an 'x' or 'y' association cannot be created, the explicit endpoint cannot be removed.

Parameters *endpoint-name* — Specifies the endpoint name.
create — Mandatory keyword to create this entry.

revert-time

Syntax **revert-time** {*revert-time* | **infinite**}
no revert-time

Context config>mirror>mirror-dest>endpoint

Description This command has an effect only when used in conjunction with a endpoint which contains a SDP of type 'primary'. It is ignored and has no effect in all other cases. The revert-timer is the delay in seconds the system waits before it switches the path of the mirror service from an active secondary SDP in the endpoint into the endpoint primary SDP after the latter comes back up.

The **no** form of the command resets the timer to the default value of 0. This means that the mirror-service path will be switched back to the endpoint primary sdp immediately after it comes back up.

Default 0 — The VLL path will be switched back to the endpoint primary SDP immediately after it comes back up.

Parameters	<i>revert-time</i> — Specifies a delay, in seconds, the system waits before it switches the path of the mirror service from an active secondary SDP in the endpoint into the endpoint primary SDP after the latter comes back up. Values 0 to 600 infinite — Forces the mirror/LI service path to never revert to the primary SDP as long as the currently active secondary -SDP is UP.

fc

Syntax	fc <i>fc-name</i> no fc
Context	config>mirror>mirror-dest
Description	This command specifies a forwarding class for all mirrored packets transmitted to the destination SAP or SDP overriding the default (be) forwarding class. All packets are sent with the same class of service to minimize out-of-sequence issues. The mirrored packet does not inherit the forwarding class of the original packet. When the destination is on a SAP, a single egress queue is created that pulls buffers from the buffer pool associated with the <i>fc-name</i>. When the destination is on an SDP, the <i>fc-name</i> defines the DiffServ-based egress queue that will be used to reach the destination. The <i>fc-name</i> also defines the encoded forwarding class of the encapsulation. The fc configuration also affects how mirrored packets are treated at the ingress queuing point on the line cards. One ingress queue is used per mirror destination (service) and that will be an expedited queue if the configured FC is expedited (one of nc, h1, ef or h2). The ingress mirror queues have no CIR, but a line-rate PIR. The no form of the command reverts 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.
Parameters	<i>fc-name</i> — 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 <i>fc-name</i> does not exist, an error will be returned and the fc command will have no effect. If the <i>fc-name</i> does exist, the forwarding class associated with <i>fc-name</i> will override the default forwarding class. Values be, l2, af, l1, h2, ef, h1, nc

isa-aa-group

Syntax	isa-aa-group <i>aa-group-id traffic-direction</i>

Context	config>mirror>mirror-dest
Description	This command specifies ISA AA group parameters.
Parameters	<i>aa-group-id</i> — Specifies the particular application group to match against to resolve to an AQP action. If set to an empty string, no match on application group is done. <i>traffic-direction</i> — Specifies the traffic directions to match against to resolve to an AQP action. This allows different policer bandwidths to apply in each direction.

mirror-dest

Syntax	mirror-dest <i>service-id</i> [type <i>mirror-type</i>] [create] no mirror-dest
Context	config>mirror
Description	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 router) or remotely, over the core of the network and have a far-end decode 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 <i>service-id</i> will receive mirrored packets from far-end router over the network core. The mirror-dest service IDs are persistent between boots of the router and are included in the configuration saves. The local sources of mirrored packets for the service ID are defined within the debug mirror mirror-source command that references the same <i>service-id</i>. Up to 255 mirror-dest service IDs can be created within a single system. The mirror-dest command is used to create or edit a service ID for mirroring purposes. If the <i>service-id</i> 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 <i>service-id</i> exists within the context of defined mirror-dest services, the CLI context is changed for editing parameters on that service ID. If the <i>service-id</i> exists within the context of another service type, an error message is returned and CLI context is not changed from the current context. LI source configuration is saved using the li>save command. The no form of the command removes a mirror destination from the system. The mirror-source or li-source associations with the mirror-dest <i>service-id</i> do not need to be removed or shutdown first. The mirror-dest <i>service-id</i> must be shutdown before the service ID can be removed. When the service ID is removed, all mirror-source or li-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 identification 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 router that this particular service is defined on.

If particular a 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.

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

Values

<i>service-id:</i>	1 to 2147483647
<i>svc-name:</i>	64 characters maximum

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

Values The following values apply to the 7750 SR:
ether, frame-relay, ppp, ip-only, atm-sdu, satop-e1, satop-e3, satop-t1, cesopsn, cesopsn-cas

Values The following values apply to the 7950 XRS:
ether, ip-only

Values The following values apply to the 7450 ESS:
ether, frame-relay, ppp

remote-source

Syntax [no] **remote-source**

Context config>mirror>mirror-dest

Description This command is used on a destination router in a remote mirroring solution. The mirroring (packet copy) is performed on the source router and sent via an SDP to the destination router. Remote mirroring requires remote-source configuration on the destination router.

Remote mirroring allows a destination router to terminate SDPs from multiple remote source routers. This allows consolidation of packet sniffers/analyzers at a single or small set of points in a network (e.g., a sniffer/analyze farm, or lawful interception gateway).

A remote-source entry must be configured on the destination router for each source router from which mirrored traffic is being sent via SDPs.

A mirror destination service that is configured for a destination router must not be configured as for a source router.

Remote-source configuration is not applicable when routable LI encapsulation is being used on the mirror source router. Remote-source configuration is only used when a source router is sending mirrored traffic to a destination router via SDPs.

Two types of remote-source entries can be configured:

- far-end
- spoke-sdp

Certain remote-source types are applicable with certain SDP types. For descriptions of the command usage in the mirror-dest context, see [far-end](#) and [spoke-sdp](#).

The 'no' form of the command removes all remote-source entries.

Default No remote source devices defined

far-end

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

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

Description This command is used on a destination router in a remote mirroring solution. See the description of the remote-source command for additional information.

When using L2TPv3, MPLS-TP or LDP IPv6 LSP SDPs in the remote mirroring solution, the destination node should be configured with **remote-src>spoke-sdp** entries. For all other types of SDPs, **remote-source>far-end** entries are used.

Up to 50 far-end entries can be specified.

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 to 223.255.255.254

vc-id *vc-id* — This is the virtual circuit identifier of the remote source. For mirror services, the *vc-id* defaults to the *service-id*. However, if the *vc-id* is being used by another service a unique *vc-id* is required to create an SDP binding. For this purpose the mirror service SDP bindings accept *vc-ids*. This VC ID must match the VC ID used on the spoke-sdp that is configured on the source router.

ing-svc-label *ing-svc-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 egress service label being used on the spoke-sdp that is configured on the source router. It must be within the range specified for manually configured service labels defined on this router. It may be reused for other far end addresses on this *mirror-dest-service-id*.

Values 2048 to 18431

tldp — Specifies that the label is obtained through signaling via the LDP.

icb — Specifies that the remote source is an inter-chassis backup SDP binding.

sap

Syntax	sap <i>sap-id</i> [create] [no-endpoint] sap <i>sap-id</i> [create] endpoint <i>name</i> no sap
Context	config>mirror>mirror-dest
Description	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 may be defined on an Ethernet access port with a dot1q, null, or q-in-q encapsulation type. 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 interface. If the interface is defined as network, the SAP creation returns an error. 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	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. endpoint <i>name</i> — 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.

cem

Syntax	cem
Context	config>mirror>mirror-dest>sap
Description	This command enables the context to specify circuit emulation (CEM) mirroring properties. Ingress and egress options cannot be supported at the same time on a CEM encap-type SAP. The options must be configured in either the ingress or egress contexts.

packet

Syntax	packet jitter-buffer <i>milliseconds</i> [payload-size <i>bytes</i>] packet <i>payload-size</i> <i>bytes</i> no packet <i>bytes</i>
Context	config>mirror>mirror-dest>sap>cem
Description	This command specifies the jitter buffer size, in milliseconds, and payload size, in bytes.
Default	The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots:

Endpoint Type	Timeslots	Default Jitter Buffer (in ms)
unstructuredE1	n/a	5
unstructuredT1	n/a	5
unstructuredE3	n/a	5
unstructuredT3	n/a	5
nxDS0 (E1/T1)	N = 1	32
	N = 2 to 4	16
	N = 5 to 15	8
	N >= 16	5
nxDS0WithCas (E1)	N	8
nxDS0WithCas (T1)	N	12

Parameters	<i>milliseconds</i> — Specifies the jitter buffer size in milliseconds (ms). Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffers is not allowed. Setting the jitter butter value to 0 sets it back to the default value. Values 1 — 250

payload-size bytes — Specifies the payload size (in bytes) of packets transmitted to the packet service network (PSN) by the CEM SAP. This determines the size of the data that will be transmitted over the service. If the size of the data received is not consistent with the payload size then the packet is considered malformed.

Default The default value depends on the CEM SAP endpoint type, and if applicable, the number of timeslots:

Endpoint Type	Timeslots	Default Payload Size (in bytes)
unstructuredE1	n/a	256
unstructuredT1	n/a	192
unstructuredE3	n/a	1024
unstructuredT3	n/a	1024
nxDS0 (E1/T1)	N = 1	64
	N = 2 to 4	N x 32
	N = 5 to 15	N x 16
	N >= 16	N x 8
nxDS0WithCas (E1)	N	N x 16
nxDS0WithCas (T1)	N	N x 24

For all endpoint types except for nxDS0WithCas, the valid payload size range is from the default to 2048 bytes.

For nxDS0WithCas, the payload size divide by the number of timeslots must be an integer factor of the number of frames per trunk multiframe (for example, 16 for E1 trunk and 24 for T1 trunk).

For 1xDS0, the payload size must be a multiple of 2.

For NxDS0, where N > 1, the payload size must be a multiple of the number of timeslots.

For unstructuredE1, unstructuredT1, unstructuredE3 and unstructuredT3, the payload size must be a multiple of 32 bytes.

Configuring the payload size and jitter buffer to values that result in less than 2 packet buffers or greater than 32 packet buffer is not allowed.

Setting the payload size to 0 sets it back to the default value.

Values 0, 16 to 2048

rtp-header

Syntax	[no] rtp-header
Context	config>mirror>mirror-dest>sap>cem
Description	This command specifies whether an RTP header is used when packets are transmitted to the packet service network (PSN) by the CEM SAP.

Default no rtp-header

egress

Syntax **egress**

Context config>mirror>mirror-dest>sap

Description This command enables access to the context to associate an egress SAP Quality of Service (QoS) policy with a mirror destination SAP.

 If no QoS policy is defined, the system default SAP egress QoS policy is used for egress processing.

ip-mirror

Syntax **ip-mirror**

Context config>mirror>mirror-dest>sap>egress

Description This command configures IP mirror information.

sa-mac

Syntax **sa-mac** *ieee-address* **da-mac** *ieee-address*
 no sa-mac

Context config>mirror>mirror-dest>sap>egress>ip-mirror

Description This command configures the source and destination MAC addresses for IP mirroring.

Parameters **sa-mac** *ieee-address* — Specifies the source MAC address. Multicast, Broadcast and zeros are not allowed.

da-mac *ieee-address* — Specifies the destination MAC address. Zeros are not allowed.

qos

Syntax **qos** *policy-id*
 no qos

Context config>mirror>mirror-dest>sap>egress

Description This command associates a QoS policy with an egress SAP for a mirrored service.

By default, no specific QoS policy is associated with the SAP for egress, so the default QoS policy is used.

The **no** form of the command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

Default	QoS policy-id 1.
Parameters	<i>policy-id</i> — The QoS policy ID to associate with SAP for the mirrored service. The policy ID must already exist.
Values	1 to 65535

service-name

Syntax	service-name <i>service-name</i> no service-name
Context	config>mirror>mirror-dest
Description	This command specifies an existing 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.

slice-size

Syntax	slice-size <i>bytes</i> no slice-size
Context	config>mirror>mirror-dest
Description	This command enables mirrored frame truncation and specifies the maximum size, in bytes, of a mirrored frame that can be transmitted to the mirror destination. This command enables mirroring larger frames than the destination packet decode equipment can handle. It also allows conservation of mirroring resources by limiting the size of the packet stream through the router and the core network. When defined, the mirror slice-size creates a threshold that truncates a mirrored frame to a specific size. For example, if the value of 256 bytes is defined, a frame larger than 256 bytes will only have the first 256 bytes transmitted to the mirror destination. The original frame is not affected by the truncation. The mirrored frame size may increase if encapsulation information is added during transmission through the network core or out the mirror destination SAP to the packet/protocol decode equipment.

The actual capability of the router to transmit a sliced or non-sliced frame is also dictated by the mirror destination SDP **path-mtu** and/or the mirror destination SAP physical MTU. Packets that require a larger MTU than the mirroring destination supports are discarded if the defined **slice-size** does not truncate the packet to an acceptable size.

Notes:

- When configuring IP mirroring, packet slice will be rejected as an incorrect option as it will cause IP packets to be rejected by the next hop with an IP header verification error.
- Slice-size is not supported by CEM encap-types or IP-mirroring.

The **no** form of the command disables mirrored packet truncation.

Default	no slice-size — Mirrored packet truncation is disabled.
Parameters	<i>bytes</i> — The number of bytes to which mirrored frames will be truncated, expressed as a decimal integer.
Values	128 to 9216

spoke-sdp

Syntax	spoke-sdp <i>sdp-id:vc-id</i> [create] [no-endpoint] spoke-sdp <i>sdp-id:vc-id</i> [create] endpoint <i>name</i> [icb] no sdp <i>sdp-id:vc-id</i>
Context	config>mirror>mirror-dest config>mirror>mirror-dest>remote-source
Description	This command binds an existing (mirror) service distribution path (SDP) to the mirror destination service ID.

Spoke SDPs are used to send and receive mirrored traffic between mirror source and destination routers in a remote mirroring solution. A spoke SDP configured in the remote-source context (**remote-src>spoke-sdp**) is used on the destination router. A spoke SDP configured in the mirror service context (**mirror-dest>spoke-sdp**) is used on the source router.

The destination node should be configured with **remote-src>spoke-sdp** entries when using L2TPv3, MPLS-TP or LDP IPv6 LSP SDPs in the remote mirroring solution. For all other types of SDPs, **remote-source>far-end** entries should be used.

Spoke SDPs are not applicable when routable LI encapsulation is employed (mirror-dest>encap).

A mirror destination service that is configured for a destination router must not be configured as for a source router.

The **no** form of the command removes the SDP binding from the mirror destination service.

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	<i>sdp-id[:vc-id]</i> — 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 <i>vc-id</i> defaults to the <i>service-id</i>. However, there are scenarios where the <i>vc-id</i> 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 <i>vc-ids</i>. Values 1 to 17407 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. icb — Indicates that the SDP is of type Inter-Chassis Backup (ICB). This is a special pseudowire used for MC-LAG and pseudowire redundancy application. An explicitly named endpoint can have a maximum of one SAP and one ICB. Once a SAP is added to the endpoint, only one more object of type ICB SDP is allowed. The ICB SDP cannot be added to the endpoint if the SAP is not part of a MC-LAG instance. This means that all other SAP types cannot exist on the same endpoint as an ICB SDP since non Ethernet SAP cannot be part of a MC-LAG instance. Conversely, a SAP which is not part of a MC-LAG instance cannot be added to an endpoint which already has an ICB SDP. An explicitly named endpoint, which does not have a SAP object, can have a maximum of four SDPs, which can include any of the following: a single primary SDP, one or many secondary SDPs with precedence, and a single ICB SDP. Default Null. The user should explicitly configure this option at create time. The user can remove the ICB type simply by retyping the SDP configuration without the icb keyword.

control-channel-status

Syntax	[no] control-channel-status
Context	<pre>config>mirror>mirror-dest>remote-src>spoke-sdp config>mirror>mirror-dest>spoke-sdp></pre>
Description	This command enables the configuration of static pseudowire status signaling on a spoke-SDP for which signaling for its SDP is set to OFF. For more information about control channel status configuration for the spoke-sdp, see the SR OS Services Guide.
Default	no control-channel-status

acknowledgment

Syntax	[no] acknowledgment
Context	config>mirror>mirror-dest>remote-src>spoke-sdp>control-channel-status config>mirror>mirror-dest>spoke-sdp>control-channel-status
Description	This command enables the acknowledgment of control channel status messages. By default, no acknowledgment packets are sent.

refresh-timer

Syntax	refresh-timer <i>value</i> no refresh-timer
Context	config>mirror>mirror-dest>remote-src>spoke-sdp>control-channel-status config>mirror>mirror-dest>spoke-sdp>control-channel-status
Description	This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.
Default	no refresh-timer
Parameters	<i>value</i> — Specifies the refresh timer value.

request-timer

Syntax	request-timer <i>timer1</i> retry-timer <i>timer2</i> timeout-multiplier <i>multiplier</i> no request-timer
Context	config>mirror>mirror-dest>remote-src>spoke-sdp>control-channel-status config>mirror>mirror-dest>spoke-sdp>control-channel-status
Description	This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.
Parameters	<i>timer1</i> — Specifies the interval at which pseudowire status messages, including a reliable delivery TLV, with the “request” bit set, are sent. Values 10 to 65535 seconds <i>retry-timer timer2</i> — Specifies the timeout interval if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries. Values 0, 3 to 60 seconds

timeout-multiplier *multiplier* — If a requesting node does not receive a valid response to a pseudowire status request within this multiplier times the retry timer, then it will assume the pseudowire is down. This parameter is optional.

Values 3 to 20 seconds

control-word

Syntax	[no] control-word
Context	config>mirror>mirror-dest>remote-src>spoke-sdp>control-channel-status config>mirror>mirror-dest>spoke-sdp>control-channel-status
Description	This command enables/disables the PW control word on spoke-sdps terminated on an IES or VPRN interface. The control word must be enabled to allow MPLS-TP OAM on the spoke-sdp It is only valid for MPLS-TP spoke-sdps when used with IES and VPRN services.
Default	no control-word

egress

Syntax	egress
Context	config>mirror>mirror-dest>spoke-sdp config>mirror>mirror-dest>remote-src>spoke-sdp
Description	This command enters the context to configure spoke SDP egress parameters.

ingress

Syntax	ingress
Context	config>mirror>mirror-dest>spoke-sdp config>mirror>mirror-dest>remote-src>spoke-sdp
Description	This command enters the context to configure spoke SDP ingress parameters.

l2tpv3

Syntax	l2tpv3
Context	config>mirror>mirror-dest>spoke-sdp>egress config>mirror>mirror-dest>remote-src>spoke-sdp>ingress

Description This command enters the context to configure an RX/TX cookie for L2TPv3 egress spoke-SDP or for the remote-source ingress spoke-sdp.

cookie

Syntax **cookie** *cookie1-value* [*cookie2-value*]
no cookie

Context config>mirror>mirror-dest>spoke-sdp>egress>l2tpv3
config>mirror>mirror-dest>remote-src>spoke-sdp>ingress>l2tpv3

Description This command configures the RX/TX cookie for L2TPv3 spoke-SDPs for the mirror destination. The command can configure L2TPv3 a single cookie for the egress spoke-SDP or one or two cookies for the remote-source ingress spoke-sdp.

The purpose of the cookie is to provide validation against misconfiguration of service endpoints, and to ensure that the right service egress is being used.

When a cookie is not configured, SR OS assumes a value of 00:00:00:00:00:00:00:00. A cookie is not mandatory. An operator may delete the egress cookie or either or both ingress cookies.

Default no cookie1 cookie2

Parameters *cookie1-value* — Specifies a 64-bit colon separated hex value.
cookie2-value — Specifies a second 64-bit colon separated hex value.

vc-label

Syntax **vc-label** *egress-vc-label*
no vc-label [*egress-vc-label*]

Context config>mirror>mirror-dest>spoke-sdp>egress
config>mirror>mirror-dest>remote-src>spoke-sdp>egress

Description This command configures the spoke-SDP egress VC label.

Parameters *egress-vc-label* — A VC egress value that indicates a specific connection.

Values 16 to 1048575

vc-label

Syntax [**no**] **vc-label** *ingress-vc-label*

Context config>mirror>mirror-dest>spoke-sdp>ingress
config>mirror>mirror-dest>remote-src>spoke-sdp>ingress

Description	This command configures the spoke-SDP ingress VC label.
Parameters	<i>vc-label</i> — A VC ingress value that indicates a specific connection.
Values	2048 to 18431

pw-path-id

Syntax	[no] pw-path-id
Context	config>mirror>mirror-dest>remote-src>spoke-sdp config>mirror>mirror-dest>spoke-sdp
Description	This command enables the context to configure an MPLS-TP Pseudowire Path Identifier for a spoke-sdp. All elements of the PW path ID must be configured in order to enable a spoke-sdp with a PW path ID. For an IES or VPRN spoke-sdp, the pw-path-id is only valid for Ethernet spoke-sdps. The pw-path-id is only configurable if all of the following is true: • The system is using network chassis mode D • SDP signaling is off • control-word is enabled (control-word is disabled by default) • the service type is epipe, vpls, cpipe, apipe, or IES/VPRN interface • mate SDP signaling is off for vc-switched services The no form of the command deletes the PW path ID.
Default	no pw-path-id

agi

Syntax	agi <i>agi</i> no agi
Context	config>mirror>mirror-dest>remote-src>spoke-sdp>pw-path-id config>mirror>mirror-dest>spoke-sdp>pw-path-id
Description	This command configures the attachment group identifier for an MPLS-TP PW.
Parameters	<i>agi</i> — Specifies the attachment group identifier.
Values	0 to 4294967295

saii-type2

Syntax	saii-type2 <i>global-id:node-id:ac-id</i> no saii-type2
Context	config>mirror>mirror-dest>remote-src>spoke-sdp>pw-path-id config>mirror>mirror-dest>spoke-sdp>pw-path-id
Description	This command configures the source individual attachment identifier (SAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the taii-type2 of the mate spoke-sdp.
Parameters	<i>global-id</i> — Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values 0 to 4294967295 <i>node-id</i> — Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values a.b.c.d or 0 to 4294967295 <i>ac-id</i> — Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value. Values 1 to 4294967295

taii-type2

Syntax	taii-type2 <i>global-id:node-id:ac-id</i> no taii-type2
Context	config>mirror>mirror-dest>remote-src>spoke-sdp>pw-path-id config>mirror>mirror-dest>spoke-sdp>pw-path-id
Description	This command configures the target individual attachment identifier (TAII) for an MPLS-TP spoke-sdp. If this is configured on a spoke-sdp for which vc-switching is also configured (for example, it is at an S-PE), then the values must match those of the saii-type2 of the mate spoke-sdp.
Parameters	<i>global-id</i> — Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values 0 to 4294967295 <i>node-id</i> — Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. Values a.b.c.d or 0 to 4294967295

ac-id — Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, then the AC ID must be set to a locally unique value.

Values 1 to 4294967295

precedence

Syntax	precedence <i>precedence-value</i> primary no precedence
Context	config>mirror>mirror-dest>spoke-sdp>egress
Description	This command indicates that the SDP is of type secondary with a specific precedence value or of type primary. The mirror/LI service always uses the primary type as the active pseudowire and only switches to a secondary pseudowire when the primary is down. The mirror service switches the path back to the primary pseudowire when it is back up. The user can configure a timer to delay reverting back to primary or to never revert back. If the active pseudowire goes down, the mirror service switches the path to a secondary sdp with the lowest precedence value. That is, secondary SDPs which are operationally up are considered in the order of their precedence value, 1 being the lowest value and 4 being the highest value. If the precedence value is the same, then the SDP with the lowest SDP ID is selected. An explicitly named endpoint can have a maximum of one SAP and one ICB. Once a SAP is added to the endpoint, only one more object of type ICB SDP is allowed. An explicitly named endpoint, which does not have a SAP object, can have a maximum of four SDPs, which can include any of the following: a single primary SDP, one or many secondary SDPs with precedence, and a single ICB SDP. An SDP is created with type secondary and with the lowest precedence value of 4.
Parameters	<i>prec-value</i> — The precedence of the SDP. Values 1 to 4 primary — A special value of the precedence which assigns the SDP the lowest precedence and enables the revertive behavior.

encap

Syntax	encap
Context	config>mirror>mirror-dest

Description This command enters the encap branch in order to configure encapsulation options for the mirrored traffic. Note that the use of encap is mutually exclusive with sap or spoke-sdp options in the same mirror-dest. Only one type of encapsulation can be specified for a single mirror-dest. Slicing and encap are mutually exclusive in the same mirror-dest.

layer-3-encap

Syntax **layer-3-encap {ip-udp-shim| ip-gre} [create]**
no layer-3-encap

Context config>mirror>mirror-dest>encap

Description This command specifies the format of the routable encapsulation to add to each copied packet. Layer-3-encap takes precedence over ethernet-encap configuration in an li-source. No changes are allowed to the layer-3-encap once a gateway is configured.

Default no layer-3-encap

Parameters **ip-udp-shim** — indicates the type of layer-3 encapsulation is an IPv4 header, UDP header and LI-Shim. Added to the mirrored packets.
ip-gre — indicates the type of layer-3 encapsulation is an IPv4 header and GRE header. Added to the mirrored packets. Only supported with mirror-dest type ip-only.

direction-bit

Syntax **direction-bit**
no direction-bit

Context config>mirror>mirror-dest>encap>layer-3-encap

Description This command is used to steal one bit from the intercept-id in the LI-Shim and use it to indicate the direction of traffic flow for an LI session. Using a direction bit may be used by a LI Mediation Gateway to distinguish between the two directions of traffic flow for an LI session when both directions share a common mirror-dest, intercept-id and session-id. If the direction bit is enabled then the Mirror Header Version (2 bit mhv) in the LI-Shim will be set to binary 01, and the next bit after the mhv is set to 0 for ingress traffic and 1 for egress traffic.

For NAT based LI, ingress means the traffic is arriving at the node from the subscriber host (applies to the 7450 ESS and 7750 SR).

No changes are allowed to the direction-bit configuration once a gateway is configured.

Default no direction-bit

router

Syntax	router <i>router-instance</i> router service-name <i>service-name</i> no router									
Context	config>mirror>mirror-dest>encap>layer-3-encap									
Description	This command specifies the routing instance into which to inject the mirrored packets. The packets will be forwarded in the routing instance based on the configurable destination IP address in the inserted IP header. If a mirror-dest is configured to inject into a VPRN service, then that VPRN service cannot be deleted. A mirror-dest with layer-3-encap will be set to operationally down if the configured destination IP address is not reachable via an interface in the routing instance or service configured for the mirror-dest. No changes are allowed to the router configuration once a gateway is configured. A service must already exist before it is specified as a router-instance. Note that vprns and ies services share the same number space for the service-id, but ies services cannot be specified as the router-instance for routable LI encap. Forwarding of routable encapsulated LI packets out an R-VPLS interface is not supported. A mirror-dest configured with routable encapsulation can be bound to a routing instance that also has an R-VPLS bound to it but the operator must ensure that the destination of the LI packets is not reachable via any R-VPLS interfaces. Any routable encapsulated LI packets that arrive at the egress of an R-VPLS interface are discarded. Parallel use of routable LI encapsulation and R-VPLS in the same routing instance is supported as long as the mirrored packets don't egress out the R-VPLS interface.									
Default	router “Base”									
Parameters	<i>router-instance</i> — Specifies the router instance. <table><tr><td>Values</td><td colspan="2"><router-name> <service-id></td></tr><tr><td><i>router-name</i></td><td>“Base” <i>name</i></td><td>(Default is Base)</td></tr><tr><td><i>service-id</i></td><td></td><td>1 to 2147483647</td></tr></table> <i>service-name</i> — Specifies the service name. Specify a character string, 64 characters maximum.	Values	<router-name> <service-id>		<i>router-name</i>	“Base” <i>name</i>	(Default is Base)	<i>service-id</i>		1 to 2147483647
Values	<router-name> <service-id>									
<i>router-name</i>	“Base” <i>name</i>	(Default is Base)								
<i>service-id</i>		1 to 2147483647								

gateway

Syntax	gateway [create] no gateway
Context	config>mirror>mirror-dest>encap>layer-3-encap

Description	Configures the parameters to send the mirrored packets to a remote destination gateway. Once a gateway is created, no changes to the layer-3-encap type, router or direction-bit are allowed.

Default	None

ip

Syntax	ip src <i>ip-address</i> dest <i>ip-address</i> no ip

Context	config>mirror>mirror-dest>encap>layer-3-encap>gateway

Description	Configures the source IPv4 address and destination IPv4 address to use in the IPv4 header part of the routable LI encapsulation.

Default	no ip

Parameters	src <i>ip-address</i> — Specifies source IP address.

Values	a.b.c.d

dest <i>ip-address</i> — Specifies destination IP address.

Values	a.b.c.d

udp

Syntax	udp src <i>udp-port</i> dest <i>udp-port</i> no udp

Context	config>mirror>mirror-dest>encap>layer-3-encap>gateway

Description	Configures the source UDP port and destination UDP port to use in the UDP header part of the routable LI encapsulation.

Default	no udp

Parameters	<i>udp-port</i> — Specifies source UDP port.

Values	1 to 65535

<i>udp-port</i> — Specifies destination UDP port.

Values	1 to 65535

2.15.2.3 Lawful Intercept Commands

li

Syntax	li
Context	config
Description	This command configures the context to configure lawful intercept (LI) parameters.

li-filter

Syntax	li-filter
Context	config>li
Description	This command enters the li-filter branch in order to create lawful intercept filter lists and entries.

li-mac-filter

Syntax	li-mac-filter filter-name [create] no li-mac-filter filter-name
Context	config>li>li-filter
Description	This command creates a Lawful Interception (LI) MAC filter list, or enters the CLI context for a LI MAC filter list. LI MAC filters are used as a manner to create confidential MAC filter based li-source entries. The LI MAC filter entries are inserted/merged into normal MAC filters as configured via the li-filter-associations and li-filter-block-reservation commands, but the LI MAC filter entries are not visible to users without LI permissions.
Parameters	<i>filter-name</i> — Specifies the name of the MAC filter. Filter names cannot start with an underscore character (e.g. “_my-filter”) and cannot use the name “default”.

li-ip-filter

Syntax	li-ip-filter filter-name [create] no li-ip-filter filter-name
Context	config>li>li-filter

Description	This command creates a Lawful Interception (LI) IPv4 filter list, or enters the CLI context for a LI IPv4 filter list. LI IPv4 filters are used as a manner to create confidential IPv4 filter based li-source entries. The LI IPv4 filter entries are inserted/merged into normal IPv4 filters as configured via the li-filter-associations and li-filter-block-reservation commands, but the LI IPv4 filter entries are not visible to users without LI permissions.
Parameters	<i>filter-name</i> — Specifies the name of the IPv4 address filter. Filter names cannot start with an underscore character (e.g. “_my-filter”) and cannot use the name “default”.

li-ipv6-filter

Syntax	li-ipv6-filter filter-name [create] no li-ipv6-filter filter-name
Context	config>li>li-filter
Description	This command creates a Lawful Interception (LI) IPv6 filter list, or enters the CLI context for a LI IPv6 filter list. LI IPv6 filters are used as a manner to create confidential IPv6 filter based li-source entries. The LI IPv6 filter entries are inserted/merged into normal IPv6 filters as configured via the li-filter-associations and li-filter-block-reservation commands, but the LI IPv6 filter entries are not visible to users without LI permissions.
Parameters	<i>filter-name</i> — Specifies the name of the IPv6 address filter. Filter names cannot start with an underscore character (e.g. “_my-filter”) and cannot use the name “default”.

entry

Syntax	entry li-entry-id [create] no entry li-entry-id
Context	config>li>li-filter>li-ip-filter config>li>li-filter>li-ipv6-filter config>li>li-filter>li-mac-filter
Description	This command creates or edits a Lawful Interception filter entry. Multiple entries can be created using unique entry-id numbers within the filter. An entry in a LI filter always has an implicit action of “forward”. The no form of the command removes the specified entry from the filter. Entries removed from the filter are immediately removed from all services or network ports where the associated filter is applied. LI filter entries can be used as li-source entries.

The entry numbers for LI filters serve purely as keys for managing the entries (deleting entries, etc). The order of LI filter entries is not guaranteed to match the entry numbers and s/w may reorder entries. Operators must use LI entries in a manner such that relative order of the LI entries amongst themselves is not important.

Parameters *li-entry-id* — Identifies the Lawful Interception filter entry.

Values 1 to 65536

match

Syntax match [frame-type {802dot3 | 802dot2-llc | 802dot2-snap | ethernet_II}]
no match

Context config>li>li-filter>li-mac-filter>entry

Description	This command creates the context for entering/editing match criteria for the filter entry and specifies an Ethernet frame type for the entry.

If more than one match criteria (within one match statement) are configured then all criteria must be satisfied (AND function) for a match to occur.

A match context may consist of multiple match criteria, but multiple match statements cannot be entered per entry.

The **no** form of the command removes the match criteria for the entry.

Parameters **frame-type** — Filters can continue to be edited by all users even when an li-source references an entry in that filter.

Values	802dot3, 802dot2-llc, 802dot2-snap, ethernet II

Default 802dot3

802dot3 — Specifies the frame type is Ethernet IEEE 802.3.

802dot2-llc — Specifies the frame type is Ethernet IEEE 802.2 LLC.

802dot2-snap — Specifies the frame type is Ethernet IEEE 802.2 SNAP.

ethernet_II — Specifies the frame type is Ethernet Type II.

match

Syntax **match** [protocols *protocols-id*]
no match

Context config>li>li-filter>li-ip-filter>entry

Description	This command enables context to enter match criteria for LI IPv4 filter and optionally allows specifying protocol value to match on.

If more than one match criterion are configured then all criteria must be satisfied for a match to occur (logical “AND”). Multiple criteria must be configured within a single match context for a given entry.

The **no** form removes the match criteria for the entry

Parameters **protocol** — The **protocol** keyword configures an IP protocol to be used as an IP filter match criterion. The protocol type such as TCP or UDP is identified by its respective protocol number.

protocol-id — Configures the decimal value representing the IP protocol to be used as an IP filter match criterion. Well known protocol numbers include ICMP(1), TCP(6), UDP(17). The **no** form the command removes the protocol from the match criteria.

Values 0 to 255 (values can be expressed in decimal, hexadecimal, or binary - DHB)

Keywords for the 7750 SR:

none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, ipv6-opts, ipv6-route, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

Keywords for the 7450 ESS:

none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp

* — udp/tcp wildcard

Protocol	Protocol ID	Description
icmp	1	Internet Control Message
igmp	2	Internet Group Management
ip	4	IP in IP (encapsulation)
tcp	6	Transmission Control
egp	8	Exterior Gateway Protocol
igp	9	Any private interior gateway (used by Cisco for IGRP)
udp	17	User Datagram
rdp	27	Reliable Data Protocol
ipv6	41	IPv6
ipv6-route	43	Routing Header for IPv6
ipv6-frag	44	Fragment Header for IPv6
idrp	45	Inter-Domain Routing Protocol
rsvp	46	Reservation Protocol
gre	47	General Routing Encapsulation
ipv6-icmp	58	ICMP for IPv6
ipv6-no-nxt	59	No Next Header for IPv6

Protocol	Protocol ID	Description
ipv6-opts	60	Destination Options for IPv6
iso-ip	80	ISO Internet Protocol
eigrp	88	EIGRP
ospf-igp	89	OSPFIGP
ether-ip	97	Ethernet-within-IP Encapsulation
encap	98	Encapsulation Header
pnni	102	PNNI over IP
pim	103	Protocol Independent Multicast
vrrp	112	Virtual Router Redundancy Protocol
l2tp	115	Layer Two Tunneling Protocol
stp	118	Spanning Tree Protocol
ptp	123	Performance Transparency Protocol
isis	124	ISIS over IPv4
crtp	126	Combat Radio Transport Protocol
crudp	127	Combat Radio User Datagram

match

Syntax	match [next-header <i>next-header</i>] no match
Context	config>li>li-filter>li-ipv6-filter>entry
Description	This command enables context to enter match criteria for LI IPv6 filter and optionally allows specifying IPv6 next-header value to match on. If more than one match criterion are configured then all criteria must be satisfied for a match to occur (logical “AND”). Multiple criteria must be configured within a single match context for a given entry. The no form removes the match criteria for the entry
Parameters	<i>next-header</i> — Specifies the IPv6 next header to match. Note that this parameter is analogous to the protocol parameter used in IP-Filter match criteria. Values [0 to 42 45 to 49 52 to 59 61 to 255] — protocol numbers accepted in decimal, hexadecimal, or binary - DHB Keywords: none, crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-icmp, ipv6-no-nxt, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, stp, tcp, udp, vrrp * — udp/tcp wildcard

dst-mac

Syntax	dst-mac <i>ieee-address</i> [<i>mask</i>] no dst-mac
Context	config>li>li-filter>li-mac-filter>entry>match
Description	Configures a destination MAC address or range to be used as a MAC filter match criterion. The no form of the command removes the destination mac address as the match criterion.
Default	no dst-mac
Parameters	<i>ieee-address</i> — Enter the 48-bit IEEE mac address to be used as a match criterion.

Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask — This 48-bit mask can be configured using:

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

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

Default 0xFFFFFFFFFFFFFFF (exact match)

Values 0x000000000000000 — 0xFFFFFFFFFFFFFFF

src-mac

Syntax	src-mac <i>ieee-address</i> [<i>ieee-address-mask</i>] no src-mac
Context	config>li>li-filter>li-mac-filter>entry>match
Description	Configures a source MAC address or range to be used as a MAC filter match criterion. The no form of the command removes the source mac as the match criteria.
Default	no src-mac
Parameters	<i>ieee-address</i> — Enter the 48-bit IEEE mac address to be used as a match criterion.
Values	HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit

ieee-address-mask — This 48-bit mask can be configured using:

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

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

Default 0xFFFFFFFFFFFFFFF (exact match)
Values 0x000000000000000 to 0xFFFFFFFFFFFFFFF

dst-ip

Syntax **dst-ip** {ip-address/mask | *ip-address ipv4-address-mask*}

Context config>li>li-filter>li-ip-filter>entry>match

Description This command configures destination IP address LI filter match criterion.

The **no** form of this command removes any configured destination IP address. The match criterion is ignored.

Default none

Parameters *ip-address* — Any address specified as dotted quad.

Values a.b.c.d

mask — Eight 16-bit hexadecimal pieces representing bit match criteria.

Values 1 to 32

ipv4-address-mask — Any mask expressed in dotted quad notation.

Values 0.0.0.0 to 255.255.255.255

dst-ip

Syntax **dst-ip** {ipv6-address/prefix-length | *ipv6-address ipv6-address-mask*}

Context config>li>li-filter>li-ip-filter>entry>match

Description This command configures destination IPv6 address LI filter match criterion.

The **no** form of this command removes any configured destination IPv6 address. The match criterion is ignored.

Default	none
Parameters	<i>ipv6-address</i> — Any IPv6 address entered as:. Values x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x - [0 to FFFF]H d - [0 to 255]D <i>prefix-length</i> — Prefix length. Values 1 to 128 <i>ipv6-address-mask</i> — Any IPv6 address mask expressed as: Values x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x - [0 to FFFF]H d - [0 to 255]D

dst-port

Syntax	dst-port {<i>lt</i> <i>gt</i> <i>eq</i>} <i>dst-port-number</i> dst-port <i>port-list-name</i> dst-port range <i>dst-port-number</i> <i>dst-port-number</i> no dst-port
Context	config>li>li-filter>li-ip-filter>entry>match config>li>li-filter>li-ipv6-filter>entry>match
Description	This command configures a destination TCP or UDP port number or port range for an IP LI filter match criterion. Note that an entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. The no form of the command removes the destination port match criterion.
Default	none
Parameters	lt gt eq — Specifies the operator to use relative to <i>dst-port-number</i> for specifying the port number match criteria. lt Specifies all port numbers less than <i>dst-port-number</i> match. gt Specifies all port numbers greater than <i>dst-port-number</i> match. eq Specifies that <i>dst-port-number</i> must be an exact match. eq — Specifies the operator to use relative to <i>dst-port-number</i> for specifying the port number match criteria. The eq keyword specifies that <i>dst-port-number</i> must be an exact match.

dst-port-number — The destination port number to be used as a match criteria expressed as a decimal integer.

Values 0 to 65535

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

range *start end* — Specifies an inclusive range of port numbers to be used as a match criteria. The destination port numbers *start-port* and *end-port* are expressed as decimal integers.

Values 0 to 65535

src-ip

Syntax	src-ip {ip-address/mask <i>ip-address ipv4-address-mask</i> }
Context	config>li>li-filter>li-ip-filter>entry>match
Description	This command configures source IP address LI filter match criterion. The no form of this command removes any configured source IP. The match criterion is ignored.
Default	no src-ip
Parameters	<i>ip-address</i> — Any address specified as dotted quad. Values a.b.c.d <i>mask</i> — Eight 16-bit hexadecimal pieces representing bit match criteria. Values 1 to 32 <i>ipv4-address-mask</i> — Any mask expressed in dotted quad notation. Values 0.0.0.0 to 255.255.255.255

src-ip

Syntax	src-ip {ipv6-address/prefix-length <i>ipv6-address ipv6-address-mask</i> } no src-ip
Context	config>li>li-filter>li-ipv6-filter>entry>match
Description	This command configures source IPv6 address LI filter match criterion. The no form of this command removes any configured source IPv6 address. The match criterion is ignored.
Default	no src-ip

Parameters	<i>ipv6-address</i> — Any IPv6 address entered as:.
Values	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x - [0 to FFFF]H d - [0 to 255]D
	<i>prefix-length</i> — Prefix length.
Values	1 to 128
	<i>ipv6-address-mask</i> — Any IPv6 address mask expressed as:
Values	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x - [0 to FFFF]H d - [0 to 255]D

src-port

Syntax	src-port { lt gt eq } <i>src-port-number</i> src-port port-list <i>port-list-name</i> src-port range <i>src-port-number src-port-number</i> no src-port
Context	config>li>li-filter>li-ip-filter>entry>match config>li>li-filter>li-ipv6-filter>entry>match
Description	This command configures a source TCP or UDP port number or port range for an IP LI filter match criterion. Note that an entry containing Layer 4 match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. The no form of the command removes the source port match criterion.
Default	no src-port
Parameters	lt gt eq — Specifies the operator to use relative to src-port-number for specifying the port number match criteria. lt Specifies all port numbers less than <i>src-port-number</i> match. gt Specifies all port numbers greater than <i>src-port-number</i> match. eq Specifies that <i>src-port-number</i> must be an exact match. <i>src-port-number</i> — The source port number to be used as a match criteria expressed as a decimal integer. Values 0 to 65535 <i>port-list-name</i> — A string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes. << R12.0 >>

range *start end* — Specifies an inclusive range of port numbers to be used as a match criteria. The source port numbers *start-port* and *end-port* are expressed as decimal integers.

Values 0 to 65535

li-filter-block-reservation

Syntax	li-filter-block-reservation
Context	config>li
Description	This command enters the li-filter-block-reservation branch in order to create lawful intercept filter reservations.

li-reserved-block

Syntax	li-reserved-block block-name [create] no li-reserved-block block-name
Context	config>li>li-filter-block-reservation
Description	This command creates or edits an LI reserved block. An LI reserved block allows an operator to define where entries from an LI filter should be inserted into a normal filter. The block reserves a configurable number of entries in the normal filter that can only be used for entries inserted from associated LI filters. The LI filter entries that get inserted into the reserved block in each normal filter are not visible to non-LI operators. The block also defines to which normal filters the reservation will be applied.
Parameters	<i>block-name</i> — Specifies the name of the MAC filter. Block names cannot start with an underscore character (e.g. “_my-filter”) and cannot use the name “default”.

start-entry

Syntax	start-entry entry-id count count no start-entry
Context	config>li>li-filter-block-reservation>li-reserved-block
Description	This command defines a block of reserved filter entries that are used to insert LI filter entries into a normal filter.
Default	no start-entry
Parameters	<i>entry-id</i> — The entry identification identifies the start of a block of reserved filter entries. Values 1 to 65536

count — This parameter identifies the number of entries in the block.

Values 1 to 8192

mac-filter

Syntax	mac-filter filter-id no mac-filter
Context	config>li>li-filter-block-reservation>li-reserved-block
Description	This command configures to which normal MAC filters the entry reservation is applied.
Parameters	<i>filter-id</i> — The filter identification identifies the normal MAC filters. Values 1—65536 <name:64 char max>

ip-filter

Syntax	ip-filter filter-name create no ip-filter
Context	config>li>li-filter-block-reservation>li-reserved-block
Description	This command configures to which normal IPv4 address filters the entry reservation is applied.
Parameters	<i>filter-id</i> — The filter identification identifies the normal IPv4 address filters. Values 1 to 65536 <name:64 char max>

ipv6-filter

Syntax	ipv6-filter filter-name create no ipv6-filter
Context	config>li>li-filter-block-reservation>li-reserved-block
Description	This command configures to which normal IPv6 address filters the entry reservation is applied.
Parameters	<i>filter-id</i> — The filter identification identifies the normal IPv6 address filters. Values 1 to 65536 <name:64 char max>

li-filter-associations

Syntax	li-filter-associations
Context	config>li
Description	This command enters the li-filter-associations branch in order to define which LI filter entries get inserted into which normal filters.

li-mac-filter

Syntax	li-mac-filter filter-name no li-mac-filter filter-name
Context	config>li>li-filter-assoc
Description	Specifies the li-mac-filter that will have its entries inserted into a list of normal mac filters.
Parameters	<i>filter-name</i> — Specifies the name of the LI MAC filter. Filter names cannot start with an underscore character (e.g. “_my-filter”) and cannot use the name “default”. 32 chars maximum.

mac-filter

Syntax	mac-filter filter-id no mac-filter filter-id
Context	config>li>li-filter-assoc>li-mac-filtr
Description	Specifies the MAC filter(s) into which the entries from the specified li-mac-filter are to be inserted. The li-mac-filter and mac-filter must already exist before the association is made. If the normal MAC filter is deleted then the association is also removed (and not re-created if the MAC filter comes into existence in the future).
Parameters	<i>filter-id</i> — The filter identification identifies the MAC filters. Values 1 to 65536 <name:64 char max>

li-ip-filter

Syntax	li-ip-filter filter-name no li-ip-filter filter-name
Context	config>li>li-filter-assoc
Description	Specifies the li-ip-filter that will have its entries inserted into a list of normal IP filters.
Parameters	<i>filter-name</i> — Specifies an existing li-ip-filter. 32 chars maximum.

ip-filter

Syntax	ip-filter <i>filter-id</i> no ip-filter <i>filter-id</i>
Context	config>li>li-filter-assoc>li-ip-fltr
Description	Specifies the IP-filter(s) into which the entries from the specified li-ip-filter are to be inserted. The li-ip-filter and ip-filter must already exist before the association is made. If the normal ip-filter is deleted then the association is also removed (and not re-created if the ip-filter comes into existence in the future).
Parameters	<i>filter-id</i> — An existing IP filter policy Values 1 to 65536 <name:64 char max>

li-ipv6-filter

Syntax	li-ipv6-filter <i>filter-name</i> no li-ipv6-filter <i>filter-name</i>
Context	config>li>li-filter-assoc
Description	Specifies the li-ipv6-filter that will have its entries inserted into a list of normal IPv6 filters.
Parameters	<i>filter-name</i> — An existing li-ipv6-filter. 32 chars maximum.

ipv6-filter

Syntax	ipv6-filter <i>filter-id</i> no ipv6-filter <i>filter-id</i>
Context	config>li>li-fltr-assoc>li-ipv6-fltr
Description	Specifies the IP-filter(s) into which the entries from the specified li-ipv6-filter are to be inserted. The li-ipv6-filter and ipv6-filter must already exist before the association is made. If the normal ipv6-filter is deleted then the association is also removed (and not re-created if the ipv6-filter comes into existence in the future).
Parameters	<i>filter-id</i> — An existing IPv6 filter policy Values 1 to 65536 <name:64 char max>

li-filter-lock-state

Syntax	li-filter-lock-state {locked unlocked-for-li-users unlocked-for-all-users} no li-filter-lock-state

Context	config>li
Description	This command configures the lock state of the filters used by LI. With the configurable filter lock for LI feature an LI user can control the behavior of filters when they are used for LI. In previous releases, when a filter entry was used as a Lawful Intercept (LI) mirror source criteria, all subsequent attempts to modify the filter were then blocked to avoid having the LI session impacted by a non-LI user. The no form of the command reverts to the default.
Default	locked
Parameters	locked — When an li-source criteria is configured that references any entry of filter Y, then filter Y can no longer be changed (until there are no longer any li-sources references to entries of filter Y). unlocked-for-li-users — Filters can continue to be edited by LI users only even when an li-source references an entry in that filter. unlocked-for-all-users — Filters can continue to be edited by all users even when an li-source references an entry in that filter.

li-source

Syntax	[no] li-source <i>service-id</i>
Context	config>li
Description	This command configures a lawful intercept (LI) mirror source.
Parameters	<i>service-id</i> — The service identification 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 router that this particular service is defined on. Values <i>service-id</i>:1 to 2147483647 <i>svc-name</i>:64 characters maximum

ip-filter

Syntax	ip-filter <i>ip-filter-id</i> [entry <i>entry-id...</i>] [intercept-id <i>intercept-id...</i>] [session-id <i>session-id...</i>] no ip-filter <i>ip-filter-id</i>
Context	config>li>li-source
Description	This command enables lawful interception (LI) 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 intercepted to the 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, IP interface or subscriber, mirroring is enabled.

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

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

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

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

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, interception 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 intercepted, no error will occur for that *entry-id* and the command will execute normally.

Parameters

ip-filter-id — The IP filter ID whose entries are to be intercepted. If the *ip-filter-id* does not exist, an error will occur and the command will not execute. Intercepting packets will commence when the *ip-filter-id* is defined on a SAP or IP interface.

entry *entry-id* — The IP filter entries to use as match criteria for lawful intercept (LI). The **entry** keyword begins a list of *entry-id*'s for interception. Multiple *entry-id* entries can be specified with a single command. Each *entry-id* must be separated by a space. Up to <N><n> 8 entry IDs may be specified in a single command.

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 criteria.

intercept-id *intercept-id* — This command configures the intercept-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This *intercept-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim**

routable encap, an *intercept-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *intercept-id* configured for an **li-source** entry, then the default value will be inserted. When the mirror service is configured with **ip-gre** routable encap, no *intercept-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4294967295 (32b) for nat li-source entries that are using a mirror service that is not configured with routable encap
1 to 1073741824 (30b) for all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and no direction-bit.
1 to 536870912 (29b) for all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and with the direction-bit enabled.

session-id *session-id* — This command configures the *session-id* that is inserted into the packet header for all mirrored packets of the associated **li-source** entry. This *session-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. The *session-id* is only valid and used for mirror services that are configured with **ip-udp-shim** routable encap (**config>mirror>mirror-dest>encap#ip-udp-shim**). For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, a *session-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *session-id* configured for an **li-source** entry, then the default value will be inserted. When a mirror service is configured with **ip-gre** routable encap, no *session-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4,294,967,295 (32b)

ipv6-filter

Syntax	ipv6-filter <i>ipv6-filter-id</i> [entry <i>entry-id...</i>] [intercept-id <i>intercept-id...</i>] [session-id <i>session-id...</i>] no ipv6-filter <i>ipv6-filter-id</i>
Context	config>li>li-source
Description	This command enables lawful interception (LI) 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 intercepted to the destination referenced by the <i>mirror-dest-service-id</i> 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 IPv6 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, IPv6 interface or subscriber, mirroring is enabled (subscriber mirroring applies only to the 7750 SR).

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

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

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

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

When the **no** command is executed with the **entry** keyword and one or more *entry-id*'s, interception 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 intercepted, no error will occur for that *entry-id* and the command will execute normally.

- Parameters**
- ipv6-filter-id* — The IPv6 filter ID whose entries are to be intercepted. If the *ipv6-filter-id* does not exist, an error will occur and the command will not execute. Intercepting packets will commence when the *ipv6-filter-id* is defined on a SAP or IPv6 interface.
 - entry** *entry-id* — The IPv6 filter entries to use as match criteria for lawful intercept (LI). The **entry** keyword begins a list of *entry-id*'s for interception. Multiple *entry-id* entries can be specified with a single command. Each *entry-id* must be separated by a space. Up to <N><n> 8 entry IDs may be specified in a single command.
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.
 - intercept-id** *intercept-id* — This command configures the intercept-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This *intercept-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim**

routable encap, an *intercept-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *intercept-id* configured for an **li-source** entry, then the default value will be inserted. When the mirror service is configured with **ip-gre** routable encap, no *intercept-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4294967295 (32b) For nat li-source entries that are using a mirror service that is not configured with routable encap

Values 1 to 1,073,741,824 (30b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and no direction-bit.

Values 1 to 536,870,912 (29b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and with the direction-bit enabled.

session-id *session-id* — This command configures the *session-id* that is inserted into the packet header for all mirrored packets of the associated **li-source** entry. This *session-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. The *session-id* is only valid and used for mirror services that are configured with **ip-udp-shim** routable encap (**config>mirror>mirror-dest>encap#ip-udp-shim**). For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, a *session-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *session-id* configured for an **li-source** entry, then the default value will be inserted. When a mirror service is configured with **ip-gre** routable encap, no *session-id* is inserted and none should be specified against the **li-source** entries.

id — The session-id value to insert into the header of the mirrored packets.

Values 1 to 4,294,967,295 (32b)

li-ip-filter

Syntax	li-ip-filter <i>filter-name</i> entry <i>li-entry-id</i> [<i>li-entry-id...</i> (upto 8 max)] [intercept-id <i>intercept-id</i> [<i>intercept-id...</i> (upto 8 max)]] [session-id <i>session-id</i> [<i>session-id...</i> (upto 8 max)]] no li-ip-filter <i>filter-name</i> [entry <i>li-entry-id</i> [<i>li-entry-id...</i> (upto 8 max)]]
Context	config>li>li-source
Description	This command enables lawful interception (LI) of packets that match specific entries in an existing LI IP filter that has been associated with a normal IP filter. The specification of an li-ip-filter entry as an li-source means that packets matching the li-ip-filter entry will be intercepted on all interfaces/saps/etc. where the associated normal ip-filter(s) are applied.
Parameters	<i>filter-name</i> — The name of the li-ip-filter. 32 characters maximum

entry *li-entry-id* — The entry id in the li-ip-filter that is to be used as an li-source criteria.

Values 1 to 65535

intercept-id *intercept-id* — This command configures the intercept-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This *intercept-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, an *intercept-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *intercept-id* configured for an **li-source** entry, then the default value will be inserted. When the mirror service is configured with **ip-gre** routable encap, no *intercept-id* is inserted and none should be specified against the **li-source** entries.

session-id *session-id* — The session-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This session-id can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. The session-id is only valid and used for mirror services that are configured with ip-udp-shim routable encap (con-fig>mirror>mirror-dest>encap#ip-udp-shim). For all types of li-source entries (filter, nat, sap, subscriber), when the mirror service is configured with ip-udp-shim routable encap, a session-id field (as part of the routable encap) is always present in the mirrored packets. If there is no session-id configured for an li-source entry, then the default value will be inserted. When a mirror service is configured with ip-gre routable encap, no session-id is inserted and none should be specified against the li-source entries.

li-ipv6-filter

Syntax	li-ipv6-filter <i>filter-name</i> entry <i>li-entry-id</i> [<i>li-entry-id...</i> (upto 8 max)] [intercept-id <i>intercept-id</i> [<i>intercept-id...</i> (upto 8 max)]] [session-id <i>session-id</i> [<i>session-id...</i> (upto 8 max)]] no li-ipv6-filter <i>filter-name</i> [entry <i>li-entry-id</i> [<i>li-entry-id...</i> (upto 8 max)]]
Context	config>li>li-source
Description	This command enables lawful interception (LI) of packets that match specific entries in an existing LI IPv6 filter that has been associated with a normal IPv6 filter. The specification of an li-ipv6-filter entry as an li-source means that packets matching the li-ipv6-filter entry will be intercepted on all interfaces/saps/etc. where the associated normal ip-filter(s) are applied.
Parameters	<i>filter-name</i> — The name of the li-ipv6-filter. 32 characters maximum. entry <i>li-entry-id</i> — The entry id in the li-ipv6-filter that is to be used as an li-source criteria. Values 1 to 65535

intercept-id *intercept-id* — The intercept-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This intercept-id can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. For all types of li-source entries (filter, nat, sap, subscriber), when the mirror service is configured with ip-udp-shim routable encap, an intercept-id field (as part of the routable encap) is always present in the mirrored packets. If there is no intercept-id configured for an li-source entry, then the default value will be inserted. When the mirror service is configured with ip-gre routable encap, no intercept-id is inserted and none should be specified against the li-source entries.

session-id *session-id* — The session-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This session-id can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. The session-id is only valid and used for mirror services that are configured with ip-udp-shim routable encap (config>mirror>mirror-dest>encap#ip-udp-shim). For all types of li-source entries (filter, nat, sap, subscriber), when the mirror service is configured with ip-udp-shim routable encap, a session-id field (as part of the routable encap) is always present in the mirrored packets. If there is no session-id configured for an li-source entry, then the default value will be inserted. When a mirror service is configured with ip-gre routable encap, no session-id is inserted and none should be specified against the li-source entries.

li-mac-filter

Syntax	li-mac-filter <i>filter-name</i> entry <i>li-entry-id</i> [<i>li-entry-id...</i> (upto 8 max)] [intercept-id <i>intercept-id</i> [<i>intercept-id...</i> (upto 8 max)]] [session-id <i>session-id</i> [<i>session-id...</i> (upto 8 max)]] no li-mac-filter <i>filter-name</i> [entry <i>li-entry-id</i> [<i>li-entry-id...</i> (upto 8 max)]]
Context	config>li>li-source
Description	This command enables lawful interception (LI) of packets that match specific entries in an existing LI MAC filter that has been associated with a normal MAC filter. The specification of an li-mac-filter entry as an li-source means that packets matching the li-mac-filter entry will be intercepted on all interfaces/saps/etc where the associated normal mac-filter(s) are applied.
Parameters	<i>filter-name</i> — The name of the li-mac-filter. 32 characters maximum. <i>li-entry-id</i> — The entry id in the li-mac-filter that is to be used as an li-source criteria. Values 1 to 65535 intercept-id <i>intercept-id</i> — This command configures the intercept-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This <i>intercept-id</i> can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. For all types of li-source entries (filter, nat, sap, subscriber), when the mirror service is configured with ip-udp-shim

routable encap, an *intercept-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *intercept-id* configured for an **li-source** entry, then the default value will be inserted. When the mirror service is configured with **ip-gre** routable encap, no *intercept-id* is inserted and none should be specified against the **li-source** entries.

session-id session-id — This command configures the *session-id* that is inserted into the packet header for all mirrored packets of the associated **li-source** entry. This *session-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. The *session-id* is only valid and used for mirror services that are configured with **ip-udp-shim** routable encap (**config>mirror>mirror-dest>encap#ip-udp-shim**). For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, a *session-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *session-id* configured for an **li-source** entry, then the default value will be inserted. When a mirror service is configured with **ip-gre** routable encap, no *session-id* is inserted and none should be specified against the **li-source** entries.

mac-filter

Syntax	mac-filter <i>mac-filter-id</i> entry [<i>entry-id...</i>] [intercept-id <i>intercept-id...</i>] [session-id <i>session-id...</i>] no mac-filter <i>mac-filter-id</i>
Context	config>li>li-source
Description	This command enables lawful interception (LI) of packets that match specific entries in an existing MAC filter. Multiple entries can be created using unique entry-id numbers within the filter. The router implementation exits the filter on the first match found and executes the actions in accordance with the accompanying action command. For this reason, entries must be sequenced correctly from most to least explicit. An entry may not have any match criteria defined (in which case, everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive. An <i>entry-id</i> within an MAC filter can only be intercepted to a single destination. If the same <i>entry-id</i> is defined multiple times, an error occurs and only the first definition is in effect. The no form of the command removes the specified entry from the IP or MAC filter. Entries removed from the IP or MAC filter are immediately removed from all services or network ports where that filter is applied.
Parameters	<i>mac-filter-id</i> — Specifies the MAC filter ID. If the <i>mac-filter-id</i> does not exist, an error will occur and the command will not execute. entry <i>entry-id</i> — The MAC filter entries to use as match criteria.

intercept-id *intercept-id* — This command configures the *intercept-id* that is inserted into the packet header for all mirrored packets of the associated **li-source** entry. This *intercept-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, an *intercept-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *intercept-id* configured for an **li-source** entry, then the default value will be inserted. When the mirror service is configured with **ip-gre** routable encap, no *intercept-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4294967295 (32b) For nat li-source entries that are using a mirror service that is not configured with routable encap

Values 1 to 1,073,741,824 (30b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and no direction-bit.

Values 1 to 536,870,912 (29b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and with the direction-bit enabled.

session-id *session-id* — This command configures the *session-id* that is inserted into the packet header for all mirrored packets of the associated **li-source** entry. This *session-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. The *session-id* is only valid and used for mirror services that are configured with **ip-udp-shim** routable encap (**config>mirror>mirror-dest>encap#ip-udp-shim**). For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, a *session-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *session-id* configured for an **li-source** entry, then the default value will be inserted. When a mirror service is configured with **ip-gre** routable encap, no *session-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4,294,967,295 (32b)

nat

Syntax	nat
Context	config>li>li-source
Description	This command enables the context to configure LI NAT parameters.

classic-lsn-sub

Syntax [no] **classic-lsn-sub router** *router-instance* **ip** *ip-address*

Context	config>li>li-source>nat
Description	This command configures a classic LSN subscriber sources. The no form of the command removes the parameter from the configuration.
Parameters	router <i>router-instance</i> — Specifies the router instance the pool belongs to, either by router name or service ID. Values <i>router-name</i>: “Base” “management” Default Base ip <i>ip-address</i> — Specifies the IP address in a.b.c.d format.

intercept-id

Syntax	intercept-id <i>id</i> no intercept-id
Context	config>li>li-source>nat>classic-lsn-sub config>li>li-source>nat>dslite-lsn-sub config>li>li-source>nat>ethernet-header config>li>li-source>nat>l2-aware-sub config>li>li-source>nat>nat64-lsn-sub
Description	This command configures the intercept-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This intercept-id can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. For nat mirroring (a nat li-source entry type), when the mirror service is not configured with any routable encap (for example, no ip-udp-shim or ip-gre configured under config>mirror>mirror-dest>encap), the presence of a configured intercept-id against an li-source (nat) entry will cause the insertion of the intercept-id after a configurable mac-da, mac-sa and etype (configured under li-source>nat>ethernet-header), at the front of each packet mirrored for that particular li-source entry. If there is no intercept-id configured (for a nat entry using a mirror service without routable encap), then a configurable mac-da and mac-sa are added to the front of the packets (but no intercept-id). In both cases a non-configurable etype is also added immediately before the mirrored customer packet. Note that routable encapsulation configured in the mirror-dest takes precedence over the ethernet-header configuration in the li-source nat entries. If routable encapsulation is configured, then the ethernet-header config is ignored and no mac header is added to the packet (the encap is determined by the mirror-dest in this case). For all types of li-source entries (filter, nat, sap, subscriber), when the mirror service is configured with ip-udp-shim routable encap, an intercept-id field (as part of the routable encap) is always present in the mirrored packets. If there is no intercept-id configured for an li-source entry, then the default value will be inserted. When the mirror service is configured with ip-gre routable encap, no intercept-id is inserted and none should be specified against the li-source entries.

The **no** form of the command removes the value from the configuration.

Default	no intercept-id (an id of 0, or no id)
Parameters	<i>id</i> — The intercept-id value to insert into the header of the mirrored packets.
Values	1 to 4294967295 (32b) For nat li-source entries that are using a mirror service that is not configured with routable encap
Values	1 to 1,073,741,824 (30b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and no direction-bit.
Values	1 to 536,870,912 (29b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and with the direction-bit enabled.

session-id

Syntax	session-id <i>id</i> no session-id
Context	config>li>li-source>nat>classic-lsn-sub config>li>li-source>nat>dslite-lsn-sub config>li>li-source>nat>ethernet-header config>li>li-source>nat>l2-aware-sub config>li>li-source>nat>nat64-lsn-sub
Description	This command configures the session-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This session-id can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. The session-id is only valid and used for mirror services that are configured with ip-udp-shim routable encap (config>mirror>mirror-dest>encap# ip-gre-shim). For all types of li-source entries (filter, nat, sap, subscriber), when the mirror service is configured with ip-udp-shim routable encap, a session-id field (as part of the routable encap) is always present in the mirrored packets. If there is no session-id configured for an li-source entry, then the default value will be inserted. When a mirror service is configured with ip-gre routable encap, no session-id is inserted and none should be specified against the li-source entries. The no form of the command removes the session-id from the configuration which results in the default value being used.
Default	no session-id (an id of 0, or no id)

Parameters *id* — The session-id value to insert into the header of the mirrored packets.

Values 1 to 4,294,967,295 (32b)

dslite-lsn-sub

Syntax **[no] dslite-lsn-sub router router-instance b4 ipv6-prefix**

Context config>li>li-source>nat

Description This command configures the Dual Stack Lite LSN subscriber source.

The **no** form of the command removes the value from the configuration.

Parameters **router router-instance** — Specifies the router instance the pool belongs to, either by router name or service ID.

Values *router-name*: “Base” | “management”

Default Base

b4 ipv6-prefix — Specifies the IPv6 address.

Values

ipv6-prefix:	<prefix>/<length>
prefix	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x to [0 to FFFF]H
	d t o[0 t o 255]D
<length>	[0 to 128]

ethernet-header

Syntax **ethernet-header [da ieee-address] [sa ieee-address] [etype ethertype]**
no ethernet-header

Context config>li>li-source>nat

Description This command configures the Ethernet header for the NAT sources.

The **no** form of the command removes the values from the configuration.

l2-aware-sub

Syntax **[no] l2-aware-sub sub-ident-string**

Context config>li>li-source>nat

Description	This command configures a Layer-2-Aware subscriber source. The no form of the command removes the values from the configuration.
Parameters	<i>sub-ident-string</i> — Specifies a source name.

sap

Syntax	sap <i>sap-id</i> {[ingress] [egress]} [intercept-id <i>intercept-id...</i>] [session-id <i>session-id...</i>] no sap <i>sap-id</i>
Context	config>li>li-source
Description	This command creates a service access point (SAP) within an LI configuration. The specified SAP must define a FastE, GigE, or XGigE, or XGigE access port with a dot1q, null, or q-in-q encapsulation type. The <i>intercept-id</i> parameter configures the intercept IDs that is inserted into the packet header for all mirrored packets of the associated li-source entry. The <i>session-id</i> parameter inserts the specified IDs into the packet header for all mirrored packets of the associated li-source entry. When the no form of this command is used on a SAP, the SAP with the specified port and encapsulation parameters is deleted.
Default	none
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. egress — Specifies that packets egressing the SAP should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification. ingress — Specifies that packets ingressing the SAP should be mirrored. Ingress packets are mirrored to the mirror destination prior to ingress packet modification. intercept-id <i>intercept-id</i> — This command configures the intercept-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This <i>intercept-id</i> can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs.

For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, an *intercept-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *intercept-id* configured for an **li-source** entry, then the default value will be inserted. When the mirror service is configured with **ip-gre** routable encap, no *intercept-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4294967295 (32b) For nat li-source entries that are using a mirror service that is not configured with routable encap

Values 1 to 1,073,741,824 (30b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and no direction-bit.

Values 1 to 536,870,912 (29b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and with the direction-bit enabled.

session-id *session-id* — This command configures the *session-id* that is inserted into the packet header for all mirrored packets of the associated **li-source** entry. This *session-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs.

The *session-id* is only valid and used for mirror services that are configured with **ip-udp-shim** routable encap (**config>mirror>mirror-dest>encap#ip-udp-shim**).

For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, a *session-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *session-id* configured for an **li-source** entry, then the default value will be inserted. When a mirror service is configured with **ip-gre** routable encap, no *session-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4,294,967,295 (32b)

subscriber

Syntax	subscriber <i>sub-ident-string</i> [sap <i>sap-id</i> [ip <i>ip-address</i>] [mac <i>ieee-address</i>]] sla-profile <i>sla-profile-name</i>] [fc {[be] [l2] [af] [l1] [h2] [ef] [h1] [nc]}} {[ingress] [egress]} [intercept-id <i>intercept-id...</i>] [session-id <i>session-id...</i>] no subscriber <i>sub-ident-string</i>
Context	config>li>li-source
Description	This command adds hosts of a subscriber to mirroring service.
Parameters	<i>sub-ident-string</i> — Specifies the name of the subscriber identification policy. <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition.

ip-address — The service IP address (system IP address) of the remote device sending LI traffic. If 0.0.0.0 is specified, any remote router is allowed to send to this service.

Values 1.0.0.1 to 223.255.255.254

mac *mac-address* — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof ip-mac** and **arp-populate**. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

sla-profile-name — Each host of a subscriber can use a different sla-profile. This option allows interception of only the hosts using the specified sla-profile. In some deployments sla-profiles are assigned per type of traffic. There can be, for example, a specific sla-profile for voice traffic (which could be used for all SIP-hosts).

Values 32 characters maximum.

fc — The name of the forwarding class with which to associate LI 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

ingress — Specifies information for the ingress policy.

egress — Specifies information for the egress policy.

intercept-id *intercept-id* — This command configures the intercept-id that is inserted into the packet header for all mirrored packets of the associated li-source entry. This *intercept-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs.

For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, an *intercept-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *intercept-id* configured for an **li-source** entry, then the default value will be inserted. When the mirror service is configured with **ip-gre** routable encap, no *intercept-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4294967295 (32b) For nat li-source entries that are using a mirror service that is not configured with routable encap

Values 1 to 1,073,741,824 (30b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and no direction-bit.

Values 1 to 536,870,912 (29b) For all types of li-source entries that are using a mirror service with routable ip-udp-shim encap and with the direction-bit enabled.

session-id *session-id* — This command configures the *session-id* that is inserted into the packet header for all mirrored packets of the associated **li-source** entry. This *session-id* can be used (for example by a downstream LI Gateway) to identify the particular LI session to which the packet belongs. The *session-id* is only valid and used for mirror services that are configured with **ip-udp-shim** routable encap (**config>mirror>mirror-dest>encap#ip-udp-shim**).

For all types of **li-source** entries (filter, nat, sap, subscriber), when the mirror service is configured with **ip-udp-shim** routable encap, a *session-id* field (as part of the routable encap) is always present in the mirrored packets. If there is no *session-id* configured for an **li-source** entry, then the default value will be inserted. When a mirror service is configured with **ip-gre** routable encap, no *session-id* is inserted and none should be specified against the **li-source** entries.

Values 1 to 4,294,967,295 (32b)

wlan-gw

Syntax	wlan-gw
Context	config>li>li-source
Description	This command enters the wlan-gw context under li-source to create li-source related configuration.
Default	none

dsm-subscriber

Syntax	[no] dsm-subscriber mac <i>xx:xx:xx:xx:xx:xx or XX-XX-XX-XX-XX-XX</i>
Context	config>li>li-source>wlan-gw
Description	This command configures the DSM UE source.
Default	none
Parameters	mac <i>xx:xx...</i> — Specifies the MAC address. Values mac-addr: <i>xx:xx:xx:xx:xx:xx</i> example: 00:0c:f1:99:85:b8 or <i>XX:XX:XX:XX:XX:XX</i> example: 00-0C-F1-99-85-B8

intercept-id

Syntax	intercept-id [1..4294967295] no intercept-id
Context	config>li>li-source>wlan-gw

Description	This command configures the intercept-id inserted in the packet header for all mirrored packets of the associated li-source. When the mirror-service is configured with the ip-udp-shim routable encaps, intercept-id field (as part of the routable encap) is always present in the mirrored packets. The intercept-id can be used by the LIG to identify a particular LI session to which the packet belongs.
Default	none
Parameters	1..4294967295 — Specifies the intercept ID inserted in the LI header.

session-id

Syntax	session-id [1..4294967295] no session-id
Context	config>li>li-source>wlan-gw
Description	This command configures the session-id inserted in the packet header for all mirrored packets of the associated li-source. When the mirror-service is configured with the ip-udp-shim routable encaps, session-id field (as part of the routable encap) is always present in the mirrored packets. The session-id can be used by the LIG to identify a particular LI session to which the packet belongs.
Default	none
Parameters	1..4294967295 — The session-id inserted in the LI header.

log

Syntax	log
Context	config>li
Description	This command enables the context to configure an event log for Lawful Intercept.

log-id

Syntax	[no] log-id log-id
Context	config>li>log
Description	This command configures an LI event log destination. The <i>log-id</i> is used to direct events, alarms/traps, and debug information to respective destinations.
Parameters	<i>log-id</i> — The log ID number, expressed as a decimal integer.
Values	1 to 100

filter

Syntax	filter <i>filter-id</i> no filter
Context	config>li>log>log-id
Description	This command adds an event filter policy with the log destination. The filter command is optional. If no event filter is configured, all events, alarms and traps generated by the source stream will be forwarded to the destination. An event filter policy defines (limits) the events that are forwarded to the destination configured in the log-id. The event filter policy can also be used to select the alarms and traps to be forwarded to a destination snmp-trap-group. The application of filters for debug messages is limited to application and subject only. Accounting records cannot be filtered using the filter command. Only one filter-id can be configured per log destination. The no form of the command removes the specified event filter from the <i>log-id</i>.
Default	no filter — No event filter policy is specified for a <i>log-id</i> .
Parameters	<i>filter-id</i> — The event filter policy ID is used to associate the filter with the <i>log-id</i> configuration. The event filter policy ID must already be defined in config>log>filter <i>filter-id</i>. Values 1 to 1000

from

Syntax	from {[li]} no from
Context	config>li>log>log-id
Description	This command configures a bit mask that specifies the log event source stream(s) to be forwarded to the destination specified in the log destination (memory, session, SNMP). Events from more than one source can be forwarded to the log destination.
Parameters	li — Specifies the li event stream that contains all events configured for Lawful Intercept activities. If the requester does not have access to the li context, the event stream will fail.

time-format

Syntax	time-format {local utc}
Context	config>li>log>log-id
Description	This command specifies whether the time should be displayed in local or Coordinated Universal Time (UTC) format.
Default	utc
Parameters	local — Specifies that timestamps are written in the system's local time. utc — Specifies that timestamps are written using the UTC value. This was formerly called Greenwich Mean Time (GMT) and Zulu time.

to

Syntax	to memory [size] to session to snmp [size]				
Context	config>li>log>log-id				
Description	This command enables the context to configure the destination type for the event log. The source of the data stream must be specified in the from command prior to configuring the destination with the to command. The to command cannot be modified or re-entered. If the destination or maximum size of an SNMP or memory log needs to be modified, the log ID must be removed and then re-created.				
Parameters	<i>size</i> — The size parameter indicates the number of events that can be stored into memory. <table> <tr> <td>Default</td><td>100</td></tr> <tr> <td>Values</td><td>50 to 1024</td></tr> </table>	Default	100	Values	50 to 1024
Default	100				
Values	50 to 1024				

save

Syntax	save
Context	config>li
Description	This command is required to save LI configuration parameters.

df-peer

Syntax	df-peer <i>df-peer-id</i> df2-addr <i>ip-address</i> df2-port <i>port</i> df3-addr <i>ip-address</i> df3-port <i>port</i> no df-peer <i>df-peer-id</i>
Context	config>li>mobile
Description	This command provisions a Delivery Function Peer, which includes Delivery Function2 used for IRI as well as Delivery Function3 used for CC, of a Lawful Intercept Gateway. The no form of the command removes the Delivery Function Peer information from the configuration.
Parameters	<i>df-peer-id</i> — Configures Delivery Function Peer parameters. Values 1 to 16 df2-addr <i>ip-address</i> — Specifies the Delivery Function2 address. This is the IP address of the Delivery Function where the IRI is to be sent. df2-port <i>port</i> — Specifies the DF2 port number. This is the TCP port of the Delivery Function where the IRI is to be sent. df3-addr <i>ip-address</i> — Specifies the Delivery Function3 address. This is the IP address of the Delivery Function where the CC is to be sent. df3-port <i>port</i> — Specifies the DF3 port number. This is the TCP port of the Delivery Function where the CC is to be sent.

local-interface

Syntax	local-interface <i>ip-address</i> [router <i>router-instance</i>] no local-interface
Context	config>li>mobile
Description	This command configures the source IP address used by the xGW/GGSN for Lawful Intercept (LI) interface. The no form of the command reverts to the default.
Default	no local-interface
Parameters	<i>ip-address</i> — Specifies the source IP address. 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 to [0 to FFFF]H d to [0 to 255]D

router *router-instance* — Specifies the router instance up to 32 characters in length.

target

Syntax	target <i>target-type</i> id <i>string</i> intercept <i>intercept</i> peer <i>df-peer-id</i> [liid <i>li-identifier</i>] no target <i>target-type</i> id <i>string</i>
Context	config>li>mobile
Description	This command configures a target for interception and assigns the Delivery Function Peer that receives the Intercept Related Information (IRI) and Content of Communication (CC) for this target. All IRI and CC messages for this target are sent to the newly specified DF peer, subsequent to target modifications. Only IMSI is currently supported as a target Identifier initially. Modifying the target command's parameters does not require a shutdown/no shutdown of the GW. The no form of the command de-activates a target that is being intercepted.
Parameters	<i>target-type</i> — Specifies the type of surveillance target identifier to be provisioned. Values <i>imsi</i>, <i>imei</i>, <i>msisdn</i> id <i>string</i> — uniquely identifies a target for the interception up to 15 characters in length. liid <i>li-identifier</i> — uniquely identifies the LI identifier up to 25 characters in length. intercept <i>intercept</i> — Specifies the interception type for the target. The intercept type is allowed to change from IRI to IRI+CC and from IRI+CC to IRI, Values <i>iri</i> — Intercept Related Information <i>iricc</i> —Intercept Related Information (IRI) and Content of Communication (CC) Default <i>iri</i> peer <i>df-peer-id</i> — Specifies the Delivery Function (DF) Peer associated with the target. Values 1 to 16 Default 1

x2-iri-qos

Syntax	[no] x2-iri-qos dscp { <i>dscp-value</i> <i>dscp-name</i> }
Context	config>li>mobile
Description	This command specifies the DSCP to be set for IRI (Intercept Related Information) messages sent to a LIG (Lawful Intercept Gateway). The no form of the command reverts to the default.

Default	disable
Parameters	<i>dscp-value</i> — Specifies the DSCP value.
Values	0 to 63
	<i>dscp-name</i> — Specifies the DSCP name.
Values	none 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 cp6 3

x3-cc-qos

Syntax	[no] x3-cc-qos dscp { <i>dscp-value</i> <i>dscp-name</i> }
Context	config>li>mobile
Description	This command specifies the DSCP to be set for CC (content of Communication) traffic sent to a LIG (Lawful Intercept Gateway). The no form of the command reverts to the default. Applies to Transport Protocol and ULIC-Header versions: • TCP with ULICv1 • UDP with ULICv1 • UDP with ULICv0
Default	disable
Parameters	<i>dscp-value</i> — Specifies the DSCP value.
Values	0 to 63
	<i>dscp-name</i> — Specifies the DSCP name.
Values	none 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 cp6 3

x3-transport

Syntax x3-transport {tcp | udp} ulic-header {v0 | v1}

Context	config>li>mobile
Description	This command specifies the transport option for an X3 interface, along with the ULIC Header version to be used. The same transport option is supported to all the Delivery Function (DF) peers in a service provider network. Changing the option requires a GW shutdown/no shutdown. Following are the valid combinations of Transport protocol and ULIC Header versions supported: • TCP with ULIC Header v1 • UDP with ULIC Header v1 • UDP with ULIC Header v0 The no form of the command reverts to the default.
Default	disable
Parameters	tcp — Specifies to use TCP as the transport option for an X3 interface. Default Only ULIC Header V1 is supported with this option. udp — Specifies to use UDP as the transport option for an X3 interface. Default Only ULIC Header V0 is supported with this option. ulic-header — Specifies the header option. v0 — Specifies ULIC v0 Header option. v1 — Specifies ULIC v1 Header option.

operator

Syntax	operator-id <i>op_id</i> no operator-id
Context	config>li>mobile
Description	This command is used to configure the operator identifier for an operator's deployment. The configured value is used to populate the operator-identifier field of the Network-Identifier IE. The "No" form of the command reverts to the default.
Default	op_id
Parameters	<i>op-id</i> — Specifies the operator identifier, string of up to 5 alphanumeric characters.

2.15.2.3.1 Other LI Configuration Commands

The following commands are also described in the *Basic System Configuration Guide*. Other LI commands are described in the *7450 ESS, 7750 SR, and 7950 XRS System Management Guide*.

li-local-save

Syntax	[no] li-local-save
Context	bof
Description	This command specifies whether or not lawful intercept (LI) configuration is allowed to be save to a local file. Modifying this command will not take affect until the system is rebooted.
Default	li-local-save

li-separate

Syntax	[no] li-separate
Context	bof
Description	This command specifies whether or not a non-LI user has access to lawful intercept (LI) information. When this command is enabled, a user who does not have LI access will not be allowed to access CLI or SNMP objects in the li context. Modifying this command will not take affect until the system is rebooted.

When the **no li-separate** command is set (the default mode), those who are allowed access to the **config>system>security>profile** context and user command nodes are allowed to modify the configuration of the LI parameters. In this mode, a user that has a profile allowing access to the **config>li** and/or **show>li** command contexts can enter and use the commands under those nodes.

When the **li-separate** command is configured, only users that have the LI access capabilities set in the **config>system>security>user>access li** context are allowed to access the **config>li** and/or **show>li** command contexts. A user who does not have LI access is not allowed to enter the **config>li** and **show>li** contexts even though they have a profile that allows access to these nodes. When in the **li-separate** mode, only users with **config>system>security>user>access li** set in their user account have the ability modify the setting LI parameters in either their own or others profiles and user configurations.

Default	no li-separate

access

Syntax	[no] access [ftp] [snmp] [console] [li]
Context	config>system>security>user
Description	This command grants a user permission for FTP, SNMP, console or lawful intercept (LI) access. If a user requires access to more than one application, then multiple applications can be specified in a single command. Multiple commands are treated additively. The no form of command removes access for a specific application. no access denies permission for all management access methods. To deny a single access method, enter the no form of the command followed by the method to be denied, for example, no access FTP denies FTP access.
Default	No access is granted to the user by default.
Parameters	ftp — Specifies FTP permission. snmp — Specifies SNMP permission. This keyword is only configurable in the config>system>security>user context and applies to the 7450 ESS and 7750 SR. console — Specifies console access (serial port or Telnet) permission. li — Allows user to access CLI commands in the lawful intercept (LI) context.

profile

Syntax	[no] profile user-profile-name
Context	config>system>security
Description	This command creates a context to create user profiles for CLI command tree permissions. Profiles are used to either deny or permit user console access to a hierarchical branch or to specific commands. Once the profiles are created, the user command assigns users to one or more profiles. You can define up to 16 user profiles but a maximum of 8 profiles can be assigned to a user. The <i>user-profile-name</i> can consist of up to 32 alphanumeric characters. The no form of the command deletes a user profile.
Default	user-profile default
Parameters	<i>user-profile-name</i> — The user profile name entered as a character string. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

li

Syntax	li
Context	config>system>security>profile
Description	This command enables the Lawful Intercept (LI) profile identifier.
Default	no li

2.16 Mirror Service Show and Debug Command Reference

2.16.1 Command Hierarchies

- [Show Commands](#)
- [Debug Commands](#)

2.16.1.1 Show Commands

```
show
  — debug [application]
  — mirror [service-id]
  — li
    — filter li-ip filter-name {counter | associations}
    — filter li-ip filter-name entry entry-id [counters]
    — filter li-ipv6 filter-name {counter | associations}
    — filter li-ipv6 filter-name entry entry-id [counters]
    — li-source [service-id]
    — log
      — log-id [log-id] [severity severity-level] [application application] [sequence
        from-seq [to-seq]] [count count] [router router-instance [expression]]
        [subject subject [regexp]] [ascending | descending]
      — mobile-gateway [target target-type id target-id] [df-peer peer-id] [summary]
      — status
  — service
    — active-subscribers summary
    — active-subscribers [subscriber sub-ident-string [sap sap-id sla-profile sla-profile-
      name]] [detail | mirror]
    — active-subscribers hierarchy [subscriber sub-ident-string]
    — service-using mirror
```

2.16.1.2 Debug Commands

```
debug
  — [no] mirror-source service-id
    — ingress-label label [label ...up to 8 max]
    — no ingress-label [label [label ...up to 8 max]]
    — ip-filter ip-filter-id entry entry-id [entry-id ...]
    — no ip-filter ip-filter-id [entry entry-id] [entry-id ...]
    — isa-aa-group isa-aa-group-id {all | unknown}
    — no isa-aa-group isa-aa-group-id
```

- **mac-filter** *mac-filter-id* **entry** *entry-id* [*entry-id* ...]
- **no mac-filter** *mac-filter-id* [**entry** *entry-id*...]
- **port** {*port-id* | **lag** *lag-id*} {[**egress**] [**ingress**]}
- **no port** {*port-id* | **lag** *lag-id*} [**egress**] [**ingress**]
- **sap** *sap-id* {[**egress**] [**ingress**]}
- **no sap** *sap-id* [**egress**] [**ingress**]
- **subscriber** *sub-ident-string* [**sap** *sap-id* [**ip** *ip-address*] [**mac** *ieee-address*] [**sla-profile** *sla-profile-name*] [**fc** {[**be**] [**l2**] [**af**] [**l1**] [**h2**] [**ef**] [**h1**] [**nc**]}}] {[**ingress**] [**egress**]}
- **no subscriber** *sub-ident-string*
- [**no**] **shutdown**

2.16.2 Command Descriptions

2.16.2.1 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

debug

Syntax	debug [<i>application</i>]
Context	show
Description	This command displays set debug points.
Parameters	<i>application</i> — Display which debug points have been set.
Values	Some examples of applications include service, ip, ospf, ospf3, bgp, mtrace, isis, mpls, rsvp, ldp, mirror, vrrp, system, filter, lag and oam
Output	

Sample Output

```
*A:EsrC# show debug
debug
  mirror-source 100
    subscriber "user1" ingress
    subscriber "user2" fc be h2 h1 nc egress
    subscriber "user3" ingress egress
    subscriber "user4" sap 1/1/2:1 fc af ef nc ingress
    subscriber "user5" sap 1/1/2:1 egress
    subscriber "user6" sap 1/1/2:1 fc be l2 af h2 ef nc ingress egress
    subscriber "user7" sap 1/1/2:1 ip 1.1.0.7 fc l1 h2 ingress
    subscriber "user8" sap 1/1/2:1 ip 1.1.0.8 fc af l1 h2 ef nc egress
```

```

        subscriber "user9" sap 1/1/2:1 ip 1.1.0.9 ingress egress
        subscriber "user10" sap 1/1/
2:1 mac 00:00:01:00:00:01 fc be 12 11 h1 nc ingress
        subscriber "user11" sap 1/1/
2:1 mac 00:00:01:00:00:02 fc be 11 h2 ef h1 egress
        subscriber "user12" sap 1/1/
2:1 mac 00:00:01:00:00:03 fc be ef ingress egress
        subscriber "user13" sap 1/1/
2:1 ip 1.1.0.13 mac 00:00:01:00:00:01 fc be ef h1 ingress
        subscriber "user14" sap 1/1/2:1 ip 1.1.0.14 mac 00:00:01:00:00:02 egress
        subscriber "user15" sap 1/1/
2:1 ip 1.1.0.15 mac 00:00:01:00:00:03 fc af 11 ef nc ingress egress
        subscriber "user16" sla-profile "sla1" ingress
        subscriber "user17" sla-profile "sla2" egress
        subscriber "user18" sla-profile "sla3" fc be af h2 ingress egress
        no shutdown
    exit
exit
*A:EsrC#

*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#

```

active-subscribers

Syntax	active-subscribers summary active-subscribers [subscriber <i>sub-ident-string</i> [sap <i>sap-id</i> sla-profile <i>sla-profile-name</i>]] [detail mirror] active-subscribers hierarchy [subscriber <i>sub-ident-string</i>]
Context	show>service
Description	This command displays active subscriber information.
Parameters	<i>sub-ident-string</i> — Specifies an existing subscriber identification string. <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. <i>sla-profile-name</i> — Specifies an existing SLA profile name. hierarchy — Keyword to display the subscriber hierarchy. <i>summary</i> — Keyword to display a subscriber summary.
Output	

Sample Output

```

*A:EsrC# show service active-subscribers mirror
=====
Active Subscribers
=====
Subscriber user1 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.1         00:00:01:00:00:01 Static
                  Ingress mirror:    100   12 af 11 nc
-----
SLA Profile Instance sap:lag-8:11 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
11.1.0.1        00:00:01:00:00:01 Static
                  Ingress mirror:    100   12 af 11 nc
-----
Subscriber user10 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.10        00:00:01:00:00:01 Static
                  Ingress mirror:    100   af ef h1 nc
-----
Subscriber user11 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.11        00:00:01:00:00:02 Static
                  Egress mirror:     100   12 ef h1
-----
Subscriber user12 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.12        00:00:01:00:00:03 Static
                  Ingress mirror:    100   be 12 af 11 h2 ef h1 nc
                  Egress mirror:    100   be 12 af 11 h2 ef h1 nc
-----
Subscriber user13 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.13        00:00:01:00:00:01 Static
                  Ingress mirror:    100   11 ef h1
-----

```



```

Subscriber user14 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.14        00:00:01:00:00:02 Static
                  Egress mirror:      100   12 h2 ef h1
-----
Subscriber user15 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.15        00:00:01:00:00:03 Static
                  Ingress mirror:      100   11 nc
                  Egress mirror:      100   11 nc
-----
Subscriber user16 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.16        00:00:01:00:00:01 Static
                  Ingress mirror:      100   be 12 af nc
-----
SLA Profile Instance sap:lag-8:11 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
11.1.0.16       00:00:01:00:00:01 Static
                  Ingress mirror:      100   be 12 af nc
-----
Subscriber user17 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla2
-----
IP Address      MAC Address      Origin
-----
1.1.0.17        00:00:01:00:00:01 Static
                  Egress mirror:      100   af 11 h1
-----
SLA Profile Instance sap:lag-8:11 - sla:sla2
-----
IP Address      MAC Address      Origin
-----
11.1.0.17       00:00:01:00:00:01 Static
                  Egress mirror:      100   af 11 h1
-----
Subscriber user18 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla3
-----
IP Address      MAC Address      Origin
-----
1.1.0.18        00:00:01:00:00:01 Static
                  Ingress mirror:      100   h2

```

```

Egress mirror:      100  h2
-----
SLA Profile Instance sap:lag-8:11 - sla:sla3
-----
IP Address      MAC Address      Origin
-----
11.1.0.18      00:00:01:00:00:01 Static
                  Ingress mirror:      100  h2
                  Egress mirror:      100  h2
-----
Subscriber user2 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.2        00:00:01:00:00:01 Static
                  Egress mirror:      100  be 12 af 11 h2 ef h1 nc
-----
SLA Profile Instance sap:lag-8:11 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
11.1.0.2        00:00:01:00:00:01 Static
                  Egress mirror:      100  be 12 af 11 h2 ef h1 nc
-----
Subscriber user3 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.3        00:00:01:00:00:01 Static
                  Ingress mirror:      100  be 12 af 11 h2 ef h1 nc
                  Egress mirror:      100  be 12 af 11 h2 ef h1 nc
-----
SLA Profile Instance sap:lag-8:11 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
11.1.0.3        00:00:01:00:00:01 Static
                  Ingress mirror:      100  be 12 af 11 h2 ef h1 nc
                  Egress mirror:      100  be 12 af 11 h2 ef h1 nc
-----
Subscriber user4 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.4        00:00:01:00:00:01 Static
                  Ingress mirror:      100  be 12 af 11 h2 ef h1 nc
-----
Subscriber user5 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----

```

```
1.1.0.5      00:00:01:00:00:01 Static
              Egress mirror:      100   be l2 af l1 h2 ef h1 nc
-----
Subscriber user6 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.6      00:00:01:00:00:01 Static
              Ingress mirror:      100   be af l1 h2
              Egress mirror:      100   be af l1 h2
-----
Subscriber user7 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.7      00:00:01:00:00:01 Static
              Ingress mirror:      100   be l2 af l1 h2 ef h1 nc
-----
Subscriber user8 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.8      00:00:01:00:00:01 Static
              Egress mirror:      100   be af l1 h1 nc
-----
Subscriber user9 (sub1)
-----
SLA Profile Instance sap:lag-8:1 - sla:sla1
-----
IP Address      MAC Address      Origin
-----
1.1.0.9      00:00:01:00:00:01 Static
              Ingress mirror:      100   be l2 af l1 h2 ef h1 nc
              Egress mirror:      100   be l2 af l1 h2 ef h1 nc
=====
*A:EsrC#
```

service-using

Syntax	service-using [mirror]
Context	show>service
Description	Displays mirror services. If no optional parameters are specified, all services defined on the system are displayed.
Parameters	mirror — Displays mirror services.
Output	

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#
```

Table 4 Show service-using Mirror 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.

li

- Syntax** li
- Context** show
- Description** Displays Lawful Intercept (LI) information.

li-source

- Syntax** li-source [service-id]
- Context** show>li

Description Displays Lawful Intercept mirror configuration and operation information.

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

Values 1 to 2147483647

Output

Sample Output

```
*A:sim138# show li li-source 2
=====
Mirror Service
=====
Service Id      : 2                Type           : Ether
Admin State     : Up              Oper State      : Up
Forwarding Class : be             Remote Sources: No
Slice           : 0
Destination SDP  : 1000 (100.1.1.2) Egress Label    : 4000
Signaling       : None

-----
Local Sources
-----
Admin State      : Up

- IP Filter      1                Entry 1
=====
*A:sim138#
```

filter

Syntax **li-ip** *filter-name* {**counter** | **associations**}
li-ip *filter-name* **entry** *entry-id* [**counters**]
li-ipv6 *filter-name* {**counter** | **associations**}
li-ipv6 *filter-name* **entry** *entry-id* [**counters**]

Context show>li

Description Displays Lawful Intercept mirror IPv4 or IPv6 address filter configuration and operation information.

Parameters *filter-name* — Specifies the LI filter name.

entry-id — Specifies the LI filter entry.

counters — Specifies LI filter counter information.

counters — Specifies LI filter association information.

log

Syntax	log
Context	show>li
Description	Displays Lawful Intercept event log information.

log-id

Syntax	log-id [<i>log-id</i>] [severity <i>severity-level</i>] [application <i>application</i>] [sequence <i>from-seq</i> [<i>to-seq</i>]] [count <i>count</i>] [router <i>router-instance</i> [expression]] [subject <i>subject</i> [regexp]] [ascending descending]
Context	show>li>log
Description	Displays information for specified log.
Parameters	<i>log-id</i> — Specifies the log ID. Values 1 to 100 <i>severity-level</i> — Specifies the severity level. Values cleared, indeterminate, critical, major, minor, warning <i>application</i> — Specifies the application name. Values bgp, cflowd, chassis, debug, igmp, lldp, mirror, ospf, pim, port, snmp, system, user, vrtr <i>from-seq</i> [<i>to-seq</i>] — Specifies the sequence value. Values 1 to 4294967295 <i>count</i> — Specifies the count. Values 1 to 4294967295 <i>subject</i> — Specifies a subject string to match. regexp — Specifies to use a regular expression match. ascending descending — Specifies the sort direction. <i>router-instance</i> — Specifies the router instance.

status

Syntax	status
Context	show>li
Description	Displays Lawful Intercept status information.

Output

Sample Output

```
*A:siml38# show li status
=====
Lawful Intercept Status Information
=====
LI Booted Config Status      : fail
LI Local Save Allowed       : yes
Separate LI administration   : no
Last LI Config Save Time    : N/A
Last Config Save Result     : none
Changes Since Last Save     : yes
Last LI Config Modified Time : 2008/01/11 10:24:30
=====
*A:siml38#
```

mobile-gateway

Syntax	mobile-gateway target <target-type> id <target-id>
Context	show>li
Description	Displays Lawful Intercept mirror configuration and operation information.
Parameters	<i>target-type</i> — Specifies the type of surveillance target identifier to be provisioned. <i>id target-id</i> — uniquely identifies a target for the interception up to 15 characters in length.

Output

Sample Output

```
show li mobile-gateway target imsi id 123456789099005
=====
LI Target Information
=====
Target id (imsi)           : 123456789099005
Intercept type             : iricc      Df peer                      : lliid: xyz123
-----
Number of targets : 1
=====
show li mobile-gateway target

=====
LI Target Information
=====
Target id (imsi)           : 123456789099005
Intercept type             : iricc      Df peer                      : lliid: xyz123
-----
-----
```

```

Target id (imsi)          : 123456789099007
Intercept type            : iricc      Df peer          : lliid: abc123
-----
-----

Number of targets : 2
=====
LI summary
=====
Total targets      : 10000          Total peers          : 1
Total IRI targets  : 0              Total IRI-CC targets: 10000
X3 transport type  : UDP            ULIC header: v1      Local interface      : 1
0.10.7.1
Router context     : Base           Operator-id: op_id
IRI-DSCP           : af41CC-DSCP: af41
=====

```

mirror

Syntax	mirror mirror-dest service-id
Context	show
Description	This command displays mirror configuration and operation information.
Parameters	<i>service-id</i> — Specify the mirror service ID.
Output	

Sample Output

```

A:SR7# show mirror mirror-dest 1000
=====
Mirror Service
=====
Service Id      : 1000          Type           : Ether
Admin State     : Up           Oper State      : Down
Forwarding Class : be          Remote Sources: No
Slice           : 0
Destination SAP : 1/1/1        Egr QoS Policy: 1
-----
Local Sources
-----
Admin State     : Up
- Port          : 1/1/2        Egress Ingress
=====
A:SR7#

```

```

A:ALA-123>config>mirror# show mirror mirror-dest 500
=====
Mirror Service
=====
Service Id      : 500          Type           : Ether
Admin State     : Up           Oper State      : Up

```



```

Forwarding Class : be                      Remote Sources: Yes
Destination SAP  : 1/1/2                  Egr QoS Policy: 1
-----
Remote Sources
-----
Far End          : 10.20.1.45              Ingress Label : 131070
-----
Local Sources
-----
Admin State      : Up
No Mirror Sources configured
=====
A:ALA-123>config>mirror#

A:ALA-456# show mirror mirror-dest 500
=====
Mirror Service
=====
Service Id       : 500                     Type           : Ether
Admin State      : Up                     Oper State      : Up
Forwarding Class : be                     Remote Sources  : No
Destination SDP   : 144 (10.20.1.44)       Egress Label    : 131070
Signaling        : TLDP
-----
Local Sources
-----
Admin State      : Up
No Mirror Sources configured
=====
A:ALA-456#

A:NS042650115# show mirror mirror-dest 100
=====
Mirror Service
=====
Service Id       : 100                     Type           : PPP
Admin State      : Up                     Oper State      : Up
Forwarding Class : be                     Remote Sources  : No
Slice           : 0Enable Port Id: Yes
Destination SDP   : 100 (2.2.2.2) Egress Label    : 131070
Signaling        : TLDP
-----
Local Sources
-----
Admin State      : Up
No Mirror Sources configured
=====
A:NS042650115#

*A:EsrC# show mirror mirror-dest 100
=====
Mirror Service
=====
Service Id       : 100                     Type           : Ether
Description      : Added by createMirrorDestination 100
Admin State      : Up                     Oper State      : Up

```

```

Forwarding Class : be                      Remote Sources: No
Slice           : 0
Destination SAP : 1/1/5:100              Egr QoS Policy: 1
-----
Local Sources
-----
Admin State      : Up
-Subs user1                                Ing
-Subs user2                                Egr
                                     FC be h2 h1 nc
-Subs user3                                Egr Ing
-Subs user4                                1/1/2:1 Ing
                                     FC af ef nc
-Subs user5                                1/1/2:1 Egr
-Subs user6                                1/1/2:1 Egr Ing
                                     FC be l2 af h2 ef nc
-Subs user7                                1/1/2:1 Ing
      IP 1.1.0.7                          FC l1 h2
-Subs user8                                1/1/2:1 Egr
      IP 1.1.0.8                          FC af l1 h2 ef nc
-Subs user9                                1/1/2:1 Egr Ing
      IP 1.1.0.9
-Subs user10                               1/1/2:1 Ing
      MAC 00:00:01:00:00:01 FC be l2 l1 h1 nc
-Subs user11                               1/1/2:1 Egr
      MAC 00:00:01:00:00:02 FC be l1 h2 ef h1
-Subs user12                               1/1/2:1 Egr Ing
      MAC 00:00:01:00:00:03 FC be ef
-Subs user13                               1/1/2:1 Ing
      IP 1.1.0.13 MAC 00:00:01:00:00:01 FC be ef h1
-Subs user14                               1/1/2:1 Egr
      IP 1.1.0.14 MAC 00:00:01:00:00:02
-Subs user15                               1/1/2:1 Egr Ing
      IP 1.1.0.15 MAC 00:00:01:00:00:03 FC af l1 ef nc
-Subs user16                               SLA sla1 Ing
-Subs user17                               SLA sla2 Egr
-Subs user18                               SLA sla3 Egr Ing
                                     FC be af h2
=====
A:EsrC#

```

[Table 5](#) lists and describes the mirroring output fields:

Table 5 Show mirror 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.
	Down — The mirror destination is administratively disabled.

Table 5 Show mirror Output Fields (Continued)

Label	Description (Continued)
Oper State	Up — The mirror destination is operationally enabled.
	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.
	No — A remote source is not configured.
Enable Port Id	Yes — PPP Port ID Mirroring is enabled.
	No — PPP Port ID Mirroring is disabled.
Slice	The value of the slice-size, the maximum portion of the mirrored frame that will be transmitted to the mirror destination. Any frame larger than the slice-size will be truncated to this value before transmission to the mirror destination. A value of 0 indicates that mirrored packet truncation based on slice size is disabled.
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.

2.16.2.2 Debug Commands

mirror-source

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

Context debug

Description 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 packet matches multiple mirror source entries (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 precedence:

- 1. Filter entry
- 2. Subscriber (applies to the 7750 SR and 7450 ESS)
- 3. SAP
- 4. Physical port

The precedence 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 a packet arriving on the SAP will be mirrored using the SAP mirror (and not mirrored using the Port mirror) because the SAP is more specific than the port.

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. If the operator enters **li>li-source svcId** for the first time, an LI source is created for the mirror service. 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	<i>service-id</i> — The mirror destination service ID for which match criteria will be defined. The <i>service-id</i> must already exist within the system.
Values	<i>service-id</i> : 1 to 2147483647 <i>svc-name</i> : 64 characters maximum

ingress-label

Syntax	ingress-label label [<i>label ...up to 8 max</i>] no ingress-label [<i>label</i> [<i>label ...up to 8 max</i>]]
Context	debug>mirror-source
Description	This command configures mirroring of ingress MPLS frames with a specific MPLS label to a mirror destination.

ip-filter

Syntax	ip-filter ip-filter-id entry entry-id [<i>entry-id ...</i>]

	no ip-filter <i>ip-filter-id</i> no ip-filter <i>ip-filter-id</i> entry <i>entry-id</i> [<i>entry-id</i> ...]
Context	debug>mirror-source
Description	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 <i>mirror-dest-service-id</i> 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 <i>entry-id</i> within an IP filter can only be mirrored to a single mirror destination. If the same <i>entry-id</i> 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 <i>entry-id</i>'s within the <i>ip-filter-id</i>. When the no command is executed with the entry keyword and one or more <i>entry-id</i>'s, mirroring of that list of <i>entry-id</i>'s is terminated within the <i>ip-filter-id</i>. If an <i>entry-id</i> is listed that does not exist, an error will occur and the command will not execute. If an <i>entry-id</i> is listed that is not currently being mirrored, no error will occur for that <i>entry-id</i> and the command will execute normally.
Default	IP filter mirroring is not defined.
Parameters	<i>ip-filter-id</i> — The IP filter ID whose entries are mirrored. If the <i>ip-filter-id</i> does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the <i>ip-filter-id</i> is defined on a SAP or IP interface. entry <i>entry-id</i> [<i>entry-id</i> ...] — The IP filter entries to use as match criteria for packet mirroring. The entry keyword begins a list of <i>entry-id</i>'s for mirroring. Multiple <i>entry-id</i> entries may be specified with a single command. Each <i>entry-id</i> must be separated by a space. If an <i>entry-id</i> 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 criteria.

isa-aa-group

Syntax	isa-aa-group <i>isa-aa-group-id</i> { all unknown } no isa-aa-group <i>isa-aa-group-id</i>
Context	debug>mirror-source
Description	This command configures AA ISAgrou as a mirror source for this mirror service. Traffic is mirrored after AA processing takes place on AA ISAs of the group, therefore, any packets dropped as part of that AA processing are not mirrored.
Parameters	<i>isa-aa-group-id</i> — Specifies the ISA ISA-AA group ID. Values 1 to 255 all — Specifies that all traffic after AA processing will be mirrored. unknown — Specifies that all traffic during the identification phase (may match policy entry or entries that have mirror action configured) and traffic that had been identified as unknown_tcp or unknown_udp after AA processing will be mirrored.

mac-filter

Syntax	mac-filter <i>mac-filter-id</i> entry <i>entry-id</i> [<i>entry-id</i> ...] no mac-filter <i>mac-filter-id</i> no mac-filter <i>mac-filter-id</i> entry <i>entry-id</i> [<i>entry-id</i> ...]
Context	debug>mirror-source
Description	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 <i>mirror-dest-service-id</i> 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 be 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.

If the MAC 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 a MAC 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 MAC filters are mirrored. Mirroring of MAC filter entries must be explicitly defined.

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	<i>mac-filter-id</i> — The MAC filter ID whose entries are mirrored. If the <i>mac-filter-id</i> does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the <i>mac-filter-id</i> is defined on a SAP. entry <i>entry-id</i> [<i>entry-id</i> ...] — The MAC filter entries to use as match criteria for packet mirroring. The entry keyword begins a list of <i>entry-id</i>'s for mirroring. Multiple <i>entry-id</i> entries may be specified with a single command. Each <i>entry-id</i> must be separated by a space. Up to 8 entry IDs may be specified in a single command. Each <i>entry-id</i> must exist within the <i>mac-filter-id</i>. If the <i>entry-id</i> is renumbered within the MAC filter definition, the old <i>entry-id</i> is removed from the list and the new <i>entry-id</i> will need to be manually added to the list if mirroring is still desired. If no <i>entry-id</i> entries are specified in the command, mirroring will not occur for that MAC filter ID. The command will have no effect.

port

Syntax	port { <i>port-id</i> lag <i>lag-id</i> } {[egress] [ingress]} no port { <i>port-id</i> lag <i>lag-id</i> } [egress] [ingress]
Context	debug>mirror-source
Description	This command enables mirroring of traffic ingressing or egressing a port (Ethernet port, SONET/SDH channel, TDM channel, 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 network, SONET/SDH, or TDM channel access. A network 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. If the port is a SONET/SDH interface, the *channel-id* must be specified to identify which channel is being mirrored (applies to the 7450 ESS and 7750 SR). Either a LAG port member or the LAG port can be mirrored.

The port is only referenced in the mirror source for mirroring purposes. The mirror source association does not need to be removed before deleting the card to which the port belongs. 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, label or filter entry, which will mirror based on a more specific criteria.

The encapsulation type on an access port or channel cannot be changed to Frame Relay if it is being mirrored (applies to the 7750 SR and 7450 ESS).

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 of the 7750 SR or 7950 XRS.
The following syntax applies to the 7750 SR:

<i>port-id</i>	<i>slot/mda/port</i> [<i>channel</i>]	
	eth-sat-id	esat-id/slot/port
		esat
		<i>id</i>
		keyword
		1 to 20
	pxc-id	pxc-id.sub-port
		pxc
		keyword
		<i>id</i>
		1 to 64
		<i>sub-port</i>
		a, b
	aps-id	aps-group-id[<i>channel</i>]
		aps
		keyword

bundle ID	<i>group-id</i>	1 to 64
	<i>bundle-type-slot/mda.bundle-num</i>	
	<i>bundle</i>	keyword
	<i>type</i>	ima, ppp
bgrp-id	<i>bundle-num</i>	1 to 336
	<i>bpgrp-type-bpgrp-num</i>	
	<i>bgrp</i>	keyword
	<i>type</i>	ima, ppp
ccag-id	<i>bgrp-num</i>	1 to 2000
	<i>ccag-id.path-id cc-type:cc-id</i>	
	<i>ccag</i>	keyword
	<i>id</i>	1 to 8
	<i>path-id</i>	a,b
	<i>cc-type</i>	sap-net, .net-sap
	<i>cc-id</i>	0 to 4094

The following syntax applies to the 7950 XRS:

<i>port-id</i>	<i>slot/mda/port [.channel]</i>	
eth-sat-id	<i>esat-id/slot/port</i>	
	<i>esat</i>	keyword
	<i>id</i>	1 to 20
pxc-id	<i>pxc-id.sub-port</i>	
	<i>pxc</i>	keyword
	<i>id</i>	1 to 64
	<i>sub-port</i>	a, b

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

Note: On the 7950 XRS, the XMA ID takes the place of the MDA.



Values 1 to 800

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	sap <i>sap-id</i> {[egress] [ingress]} no sap <i>sap-id</i> [egress] [ingress]
Context	debug>mirror-source
Description	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 <i>sap-id</i> 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 and egress 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. The same SAP cannot be associated with multiple mirror source definitions for egress 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. Note that the ingress and egress options cannot be supported at the same time on a CEM encap-type SAP. The options must be configured in either the ingress or egress contexts (applies to the 7750 SR and 7950 XRS). 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.
Parameters	<i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. <i>channel-id</i> — The SONET/SDH or TDM channel on the port of the SAP. A period separates the physical port from the <i>channel-id</i>. The port must be configured as an access port. This parameter applies only to the 7750 SR. egress — Specifies that packets egressing the SAP should be mirrored. Egress packets are mirrored to the mirror destination after egress packet modification. ingress — Specifies that packets ingressing the SAP should be mirrored. Ingress packets are mirrored to the mirror destination prior to ingress packet modification.

subscriber

Syntax	subscriber <i>sub-ident-string</i> [sap <i>sap-id</i> [ip <i>ip-address</i>] [mac <i>ieee-address</i>]] sla-profile <i>sla-profile-name</i>] [fc {[be] [l2] [af] [l1] [h2] [ef] [h1] [nc]}} {[ingress] [egress]}] no subscriber <i>sub-ident-string</i>
Context	debug>mirroring-source
Description	This command adds hosts of a subscriber to mirroring service.
Parameters	<i>sub-ident-string</i> — Specifies the name of the subscriber identification policy. <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. <i>ip ip-address</i> — The service IP address (system IP address) of the remote 7750 SR or 7450 ESS device sending LI traffic. Values 1.0.0.1 to 223.255.255.254 <i>mac mac-address</i> — Specify this optional parameter when defining a static host. The MAC address must be specified for anti-spoof ip-mac and arp-populate. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address. <i>sla-profile sla-profile-name</i> — Specifies the SLA profile name. Values 32 characters maximum. fc — The name of the forwarding class with which to associate LI traffic. Values be, l2, af, l1, h2, ef, h1, nc ingress — Specifies information for the ingress policy. egress — Specifies information for the egress policy.

ingress-label

Syntax	[no] ingress-label <i>label</i> [<i>label</i> ... up to 8 max] no ingress-label <i>label</i> [<i>label</i> ... up to 8 max]
Context	debug>mirror-source
Description	This command enables ingress MPLS frame mirroring based on the top-of-stack MPLS label. Multiple labels can be defined simultaneously. The ingress-label command is used to mirror ingressing MPLS frames with specific MPLS labels to a specific mirror destination. The ingress label must be at the top of the label stack and can only be mirrored to a single mirror destination. If the same label is defined with multiple mirror destinations, an error is generated and the original mirror destination remains.

The **ingress-label** mirror source overrides all other mirror source definitions. The MPLS frame is mirrored to the mirror destination as it is received on the ingress network port. The router MPLS label space is global for the system. A specific label is mirrored to the mirror destination regardless of the ingress interface.

By default, no ingress MPLS frames are mirrored. The **ingress-label** command must be executed to start mirroring on a specific MPLS label.

The **no ingress-label** command removes all label mirroring for the mirror source. To stop mirroring on specific labels, use the **no ingress-label label** form of the command. Multiple labels may be given in a single **no ingress-label** command.

Default	No ingress MPLS labels for mirroring are defined.
Parameters	<i>label</i> — The top-of-stack label received on ingress to be mirrored. A label can only be mirrored to a single mirror destination. If the label does not exist on any ingress network ports, no packets are mirrored for that label. An error will not occur. Once the label exists on a network port, ingress mirroring commences for that label. Values 0 to 1048575. The local MPLS stack may not support portions of this range.

3 OAM, SAA, and OAM-PM

3.1 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](#)
- [IP Performance Monitoring \(IP PM\)](#)
- [Ethernet Connectivity Fault Management \(ETH-CFM\)](#)
- [ETH-CFM Statistics](#)
- [ETH-CFM Packet Debug](#)
- [ETH-CFM CoS Considerations](#)
- [OAM Mapping](#)
- [Service Assurance Agent \(SAA\)](#)
- [OAM Performance Monitoring \(OAM-PM\)](#)
- [Traceroute with ICMP Tunneling In Common Applications](#)

3.2 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, VC-labels to a service and each service to a service tunnel 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 MPLS LSPs, SDPs, services and VPLS MACs within a service.

This section describes the following topics:

- [LSP Diagnostics: LSP Ping and Trace](#)
- [LSP Ping/Trace for an LSP Using a BGP IPv4 Label Route](#)
- [ECMP Considerations](#)
- [Lsp-ping and Lsp-trace over Unnumbered IP Interface](#)
- [Downstream Detailed Mapping \(DDMAP\) TLV](#)
- [Using DDMAP TLV in LSP Stitching and LSP Hierarchy](#)
- [LDP Tree Trace: End-to-End Testing of Paths in an LDP ECMP Network](#)
- [LDP ECMP Tree Building](#)
- [Periodic Path Exercising](#)
- [LSP Ping for RSVP P2MP LSP \(P2MP\)](#)
- [LSP Trace for RSVP P2MP LSP](#)
- [Tunneling of ICMP Reply Packets over MPLS LSP](#)
- [QoS Handling of Tunneled ICMP Reply Packets](#)
- [Summary of UDP Traceroute Behavior With and Without ICMP Tunneling](#)
- [SDP Diagnostics](#)
- [SDP Ping](#)
- [SDP MTU Path Discovery](#)
- [Service Diagnostics](#)

- [VPLS MAC Diagnostics](#)
- [MAC Ping](#)
- [MAC Trace](#)
- [CPE Ping](#)
- [CPE Ping for PBB Epipe](#)
- [MAC Populate](#)
- [MAC Purge](#)
- [VLL Diagnostics](#)
- [VCCV Ping](#)
- [Automated VCCV-Trace Capability for MS-Pseudowire](#)
- [IGMP Snooping Diagnostics](#)
- [MFIB Ping](#)
- [ATM Diagnostics](#)
- [MPLS-TP On-Demand OAM Commands](#)
- [MPLS-TP Pseudowires: VCCV-Ping/VCCV-Trace](#)
- [MPLS-TP LSPs: LSP-Ping/LSP Trace](#)
- [VxLAN Ping Supporting EVPN for VxLAN](#)
- [Show Commands](#)
- [BFD](#)

3.2.1 LSP Diagnostics: LSP Ping and Trace

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.

3.2.2 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 shown in [Figure 21](#).

Figure 21 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 [profile {in|out}]] [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] [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] [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 the Nokia router platform. BGP label IPv6 /128 prefixes are not currently usable as tunnels on the router platform and as such are not supported in LSP ping/trace.

3.2.3 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 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 users 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 that 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. This hash is based on:
 - a. The {incoming port, system interface address, label-stack} when the **Lsr-load-balancing** option of the incoming interface is configured to **lbl-only**. In this case the 127/8 prefix address entered in the **path-destination** option is not used to select the outgoing interface. All packets received with the same label stack will map to a single and same outgoing interface.
 - b. The {incoming port, system interface address, label-stack, SRC/DEST IP fields of the packet} when the **Lsr-load-balancing** option of the incoming interface is configured to **lbl-ip**. The SRC IP field corresponds to the value entered by the user in the **src-ip-address** option (default system IP interface address). The DEST IP field corresponds to the 127/8 prefix address entered in the **path-destination** option. In this case, the CPM code will map the packet, as well as any packet in a sub-range of the entire 127/8 range, to one of the possible outgoing interface of the FEC.
 - c. The {SRC/DEST IP fields of the packet} when the **Lsr-load-balancing** option of the incoming interface is configured to **ip-only**. The SRC IP field corresponds to the value entered by the user in the **src-ip-address** option (default system IP interface address). The DEST IP field corresponds to the 127/8 prefix address entered in the **path-destination** option. In this case, the CPM code will map the packet, as well as any packet in a sub-range of the entire 127/8 range, to one of the possible outgoing interface of the FEC.

In all above cases, 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.

Note that if the user enabled the **system-ip-load-balancing hash** option (**config>system>system-ip-load-balancing**), then the LSR hashing is modified by applying the system IP interface, with differing bit-manipulation, to the hash of packets of all three options (**lbl-only**, **lbl-ip**, **ip-only**). This system level option enhances the LSR packet distribution such that the probability of the same flow selecting the same ECMP interface index or LAG link index at two consecutive LSR nodes is minimized.

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 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).

3.2.4 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.

3.2.5 Downstream Detailed Mapping (DDMAP) TLV


The DDMAP TLV provides the same features as the existing DSMAP TLV, plus the enhancement 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. [Figure 22](#) shows the structures of these two objects as defined in RFC 6424.

Figure 22 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 as shown in [Figure 23](#).

Figure 23 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:

- 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.
- 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:

- 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.
- 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:

- 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.
 - 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.
 - 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.
- A sender node never includes the DSMAP or DDMAP TLV in an lsp-ping message.

3.2.6 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:

- 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.
- 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.
- 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.
- Full validation of an LDP FEC tunneled over an RSVP LSP using LSP trace.
- 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.

3.2.6.1 Responder Node Procedures

1. As a responder node, the router will always insert a global return code 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 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).

3.2.6.2 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 that 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 router 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.
 5. 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.

3.2.7 LDP Tree Trace: End-to-End Testing of Paths in an LDP ECMP Network

Figure 24 shows an IP/MPLS network which uses LDP ECMP for network resilience. Faults that are detected through IGP and/or LDP are corrected as soon as IGP and LDP re-converge. The impacted traffic will be forwarded on the next available ECMP path as determined by the hash routine at the node that had a link failure.

Figure 24 Network Resilience Using LDP ECMP




OSSG265

However, there are faults which the IGP/LDP control planes may not detect. These faults may be due to a corruption of the control plane state or of the data plane state in a node. Although these faults are very rare and mostly due to misconfiguration, the LDP Tree Trace OAM feature is intended to detect these “silent” data plane and control plane faults. For example, it is possible that the forwarding plane of a node has a corrupt Next Hop Label Forwarding Entry (NHLFE) and keeps forwarding packets over an ECMP path only to have the downstream node discard them. This data plane fault can only be detected by an OAM tool that can test all possible end-to-end paths between the ingress LER and the egress LER. A corruption of the NHLFE entry can also result from a corruption in the control plane at that node.

3.2.8 LDP ECMP Tree Building

When the LDP tree trace feature is enabled, the ingress LER builds the ECM 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. In order to build the ECMP tree, the router LER inserts an IP address range drawn from the 127/8 space. When received by the downstream LSR, it will use 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 router LER, it will record this information and proceed 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 will be used since the objective is to have the LSR downstream of the router LER pass this message to its downstream node along the first ECMP path.

The following figure illustrates the behavior through the following example adapted from RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*:



LSR A has two downstream LSRs, B and F, for PE2 FEC. PE1 receives an echo reply from A with the Multipath Type set to 4, with low/high IP addresses of 127.1.1.1->127.1.1.255 for downstream LSR B and 127.2.1.1->127.2.1.255 for downstream LSR F. PE1 reflects this information to LSR B. B, which has three downstream LSRs, C, D, and E, computes that 127.1.1.1->127.1.1.127 would go to C and 127.1.1.128->127.1.1.255 would go to D. B would then respond with 3 Downstream Mappings: to C, with Multipath Type 4 (127.1.1.1->127.1.1.127); to D, with Multipath Type 4 (127.1.1.127->127.1.1.255); and to E, with Multipath Type 0.

The router supports multipath type 0 and 8, and up to a maximum of 36 bytes for the multipath length and supports the LER part of the LDP ECMP tree building feature.

A user configurable parameter sets the frequency of running the tree trace capability. The minimum and default value is 60 minutes and the increment is 1 hour.

The router LER gets the list of FECs from the LDP FEC database. New FECs will be added to the discovery list at the next tree trace and not when they are learned and added into the FEC database. The maximum number of FECs to be discovered with the tree building feature is limited to 500. The user can configure FECs to exclude the use of a policy profile.

3.2.9 Periodic Path Exercising

The periodic path exercising capability of the LDP tree trace feature runs in the background to test the LDP ECMP paths discovered by the tree building 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 trace for this FEC.

The periodic LSP ping messages continuously probes an ECMP path at a user configurable rate of at least 1 message per minute. This is the minimum and default value. The increment is 1 minute. If an interface is down on a router LER, then LSP ping probes that normally go out this interface will not be sent.

The LSP ping routine updates the content of the MPLS echo request message, specifically the IP address, as soon as the LDP ECMP tree trace has output the results of a new computation for the path in question.

3.2.10 LSP Ping for RSVP P2MP LSP (P2MP)

The P2MP LSP ping complies to RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*.

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 [profile {in | out}]] [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 5 addresses can be specified in a single execution of the **p2mp-lsp-ping** command. If all 5 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 **ttl** 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.

For more information about P2MP refer to the *MPLS Guide*.

3.2.11 LSP Trace for RSVP P2MP LSP

The P2MP LSP trace complies to RFC 6425. An LSP trace can be generated by entering the following OAM command:


```
oam p2mp-lsp-trace lsp-name p2mp-instance instance-name s2l-dest-address ip-
address [fc fc-name [profile {in|out}]] [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 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.



```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Downstream Interface Address (4 or 16 octets)
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Return Code | Return Subcode | Subtlv Length
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.
.
.
.
.
.
.
.
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Downstream Detailed Mapping TLV format is derived from the Downstream Mapping (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.

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 must respond.

When a branch LSR or BUD LSR node responds to the sender of the echo request message, 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>".

Since a single egress LER address, for example an S2L path, can be traced, the branch LSR or bud LSR node will set the multipath type of zero in the downstream mapping TLV in the echo response message as no egress LER address need to be included.


3.2.11.1 LSP Trace Behavior When S2L Path Traverses a Re-Merge Node

When a 7450 ESS, 7750 SR or 7950 XRS LSR performs a re-merge of one or more ILMs of the P2MP LSP to which the traced S2L sub-LSP belongs, it may block the ILM over which the traced S2L resides. This causes the trace to either fail or to succeed with a missing hop.

The following is an example of this behavior.

S2L1 and S2L2 use ILMs which re-merge at node B. Depending of which ILM is blocked at B, the TTL=2 probe will either yield two responses or will timeout.

```
S2L1 = ACBDF (to leaf F)
S2L2 = ABDE (to leaf E)
```



- 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.

- 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.

3.2.12 Tunneling of ICMP Reply Packets over MPLS LSP

This feature enables the tunneling of ICMP reply packets over MPLS LSP at an LSR node as per RFC 3032. At an LSR node, including an ABR, ASBR, or data path Router Reflector (RR) node, the user enables the ICMP tunneling feature globally on the system using the **config>router>icmp-tunneling** command.

This feature supports tunneling ICMP replies to a UDP traceroute message. It does not support tunneling replies to an ICMP ping message. The LSR part of this feature consists of crafting the reply ICMP packet of type=11- 'time exceeded', with a source address set to a local address of the LSR node, and appending the IP header and leading payload octets of the original datagram. The system skips the lookup of the source address of the sender of the label TTL expiry packet, which becomes the destination address of the ICMP reply packet. Instead, CPM injects the ICMP reply packet in the forward direction of the MPLS LSP the label TTL expiry packet was received from. The TTL of pushed labels should be set to 255.

The source address of the ICMP reply packet is determined as follows:

1. The LSR uses the address of the outgoing interface for the MPLS LSP. Note that with LDP LSP or BGP LSP, multiple ECMP next-hops can exist in which case the first outgoing interface is selected.
2. If the interface does not have an address of the same family (IPv4 or IPv6) as the ICMP packet, then the system address of the same family is selected. If one is not configured, the packet is dropped.

When the packet is received by the egress LER, it performs a regular user packet lookup in the data path in the GRT context for BGP shortcut, 6PE, and BGP label route prefixes, or in VPRN context for VPRN and 6VPE prefixes. It then forwards it to the destination, which is the sender of the original packet which TTL expired at the LSR.

If the egress LER does not have a route to the destination of the ICMP packet, it drops the packets.

The rate of the tunneled ICMP replies at the LSR can be directly or indirectly controlled by the existing IOM level and CPM levels mechanisms. Specifically, the rate of the incoming UDP traceroute packets received with a label stack can be controlled at ingress IOM using the distributed CPU protection feature. The rate of the ICMP replies by CPM can also be directly controlled by configuring a system wide rate limit for packets ICMP replies to MPLS expired packets which are successfully forwarded to CPM using the command 'configure system security vprn-network-exceptions'. Note that while this command's name refers to VPRN service, this feature rate limits ICMP replies for packets received with any label stack, including VPRN and shortcuts.

The 7450 ESS, 7750 SR and 7950 XRS router implementation supports appending to the ICMP reply of type Time Exceeded the MPLS label stack object defined in RFC 4950. It does not include it in the ICMP reply type of Destination unreachable.

The new MPLS Label Stack object permits an LSR to include label stack information including label value, EXP, and TTL field values, from the encapsulation header of the packet that expired at the LSR node. The ICMP message continues to include the IP header and leading payload octets of the original datagram.

In order to include the MPLS Label Stack object, the SR OS implementation adds support of RFC 4884 which defines extensions for a multi-part ICMPv4/v6 message of type Time Exceeded. Section 5 of RFC 4884 defines backward compatibility of the new ICMP message with extension header with prior standard and proprietary extension headers.

In order to guarantee interoperability with third party implementations deployed in customer networks, the router implementation is able to parse in the receive side all possible encapsulations formats as defined in Section 5 of RFC 4884. Specifically:

The new MPLS Label Stack object permits an LSR to include label stack information including label value, EXP, and TTL field values, from the encapsulation header of the packet that expired at the LSR node. The ICMP message continues to include the IP header and leading payload octets of the original datagram.

1. If the length attribute is zero, it is treated as a compliant message and the router implementation will process the original datagram field of size equal to 128 bytes and with no extension header.
2. If the length attribute is not included, it is treated as a non-compliant message and the router implementation will process the original datagram field of size equal to 128 bytes and also look for a valid extension header following the 128 byte original datagram field. If the extension is valid, it is processed accordingly, if not it is assumed the remainder of the packet is still part of the original datagram field and process it accordingly. Note that the router implementation only validates the ICMP extension version number and not the checksum field in the extension header. The checksum of the main time exceeded message is also not validated as per prior implementation.
3. An ICMP reply message will be dropped if it includes more than one MPLS label object. In general when a packet is dropped due to an error in the packet header or structure, the traceroute will timeout and will not display an error message.
4. When processing the received ICMP reply packet, an unsupported extension header will be skipped.

In the transmit side, when the MPLS Label Stack object is added as an extension to the ICMP reply message, it is appended to the message immediately following the "original datagram" field taken from the payload of the received traceroute packet. The size of the appended "original datagram" field contains exactly 128 octets. If the original datagram did not contain 128 octets, the "original datagram" field is zero padded to 128 octets.

For sample output of the traceroute OAM tool when the ICMP tunneling feature is enabled see, [Traceroute with ICMP Tunneling In Common Applications](#).

3.2.13 QoS Handling of Tunneled ICMP Reply Packets

When the ICMP reply packet is generated in CPM, its FC is set by default to NC1 with the corresponding default ToS byte value of 0xC0. The DSCP value can be changed by configuring a different value for an ICMP application under the **config>router>sgt-qos icmp** context.

When the packet is forwarded to the outgoing interface, the packet is queued in the egress network queue corresponding to its CPM assigned FC and profile parameter values. The marking of the packet's EXP is dictated by the {FC, profile}-to-EXP mapping in the network QoS policy configured on the outgoing network interface. The TOS byte, and DSCP value for that matter, assigned by CPM are not modified by the IOM.

3.2.14 Summary of UDP Traceroute Behavior With and Without ICMP Tunneling

At a high level, the major difference in the behavior of the UDP traceroute when ICMP tunneling is enabled at an LSR node is that the LSR node tunnels the ICMP reply packet towards the egress of the LSP without looking up the traceroute sender's address. When ICMP tunneling is disabled, the LSR looks it up and replies if the sender is reachable. However there are additional differences in the two behaviors and they are summarized in the following.

- icmp-tunneling disabled/IPv4 LSP/IPv4 traceroute:
 - Ingress LER, egress LER, and LSR attempt to reply to the UDP traceroute of both IPv4 and VPN-IPv4 routes.

- For VPN-IPv4 routes, the LSR will attempt to reply but it may not find a route and in such a case the sender node will timeout. In addition, the ingress and egress ASBR nodes in VRPN inter-AS option B will not respond as in current implementation and the sender will timeout.
- icmp-tunneling disabled/IPv4 LSP/IPv6 traceroute:
 - Ingress LER and egress LER reply to traceroute of both IPv6 and VPN-IPv6 routes. LSR does not reply.
- icmp-tunneling enabled/IPv4 LSP/IPv4 traceroute:
 - ingress LER and egress LER reply directly to the UDP traceoute of both IPv4 and VPN-IPv4 routes. LSR tunnels the reply to endpoint of the LSP to be forwarded from there to the source of the traceroute.
 - For VPN-IPv4 routes, the ingress and egress ASBR nodes in VPRN inter-AS option B will also tunnel the reply to the endpoint of the LSP and as such there is no timeout at the sender node like in the case when icmp-tunneling is disabled.
- icmp-tunneling enabled/IPv4 LSP/IPv6 traceroute:
 - ingress LER and egress LER reply directly to the UDP traceoute of both IPv6 and VPN-IPv6 routes. LSR tunnels the reply to endpoint of the LSP to be forwarded from there to the source of the traceroute.
 - For VPN-IPv6 routes, the ingress and egress ASBR nodes in VPRN inter-AS option B will also tunnel the reply to the endpoint of the LSP like in the case when icmp-tunneling is disabled.

In the presence of ECMP, CPM generated UDP traceroute packets are not sprayed over multiple ECMP next-hops. The first outgoing interface is selected. In addition, a LSR ICMP reply to a UDP traceroute will also be forwarded over the first outgoing interface regardless if ICMP tunneling is enabled or not. When ICMP tunneling is enabled, it means the packet is tunneled over the first downstream interface for the LSP when multiple next-hops exist (LDP FEC or BGP label route). In all cases, the ICMP reply packet uses the outgoing interface address as the source address of the reply packet.

3.2.15 SDP Diagnostics

The router SDP diagnostics are SDP ping and SDP MTU path discovery.

3.2.16 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 router. 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

3.2.17 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).

3.2.18 Service Diagnostics

Nokia'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 router to verify round-trip connectivity and delay to the far-end of the service. Nokia's implementation functions for both GRE and 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

3.2.19 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. Nokia 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.

3.2.20 MAC Ping

For a MAC ping test, the destination MAC address (unicast or multicast) to be tested must be specified. A MAC ping packet is sent through the data plane. The ping packet goes out with the data plane format.

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.

3.2.21 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 is sent via the data plane.

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).

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.

3.2.22 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 router. 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.

3.2.23 CPE Ping for PBB Epipe

CPE ping has been supported for VPLS services since Release 3.0 of SR OS. It enables the connectivity of the access circuit between a VPLS PE and a CPE to be tested, even if the CPE is unmanaged, and therefore the service provider cannot run standardized Ethernet OAM to the CPE. The command “cpe-ping” for a specific destination IP address within a VPLS is translated into a MAC-ping towards a broadcast MAC address. All destinations within the VPLS context are reached by this ping to the broadcast the MAC address. At all these destinations, an ARP will be triggered for the specific IP address (with the IP destination address equals to the address from the request, mac-da equals to all1's, mac-sa equals to the CPM-mac-address and the IP source address, which is the address found in the request). The destination receiving a response will reply back to the requester.

Release 10.0 extended the CPE ping command for local, distributed, and PBB Epipe services provisioned over a PBB VPLS. CPE ping for Epipe implements an alternative behavior to CPE ping for VPLS that enables fate sharing of the CPE ping request with the Epipe service. Any PE within the epipe service (the source PE) can launch the CPE ping. The source PE builds an arp request and encapsulates it to be sent in the epipe as if it came from a customer device by using its chassis MAC as the source MAC address. The ARP request then egresses the remote PE device as any other packets on the epipe. The remote CPE device responds to the ARP and the reply is transparently sent on the epipe towards the source PE. The source PE will then look for a match on its chassis MAC in the inner customer DA. If a match is found, the source PE device intercepts this response packet.

This method is supported regardless of whether the network uses SDPs or SAPs. It is configured using the existing **oam>cpe-ping** CLI command.



Note: This feature does not support IPv6 CPEs.

3.2.23.1 Hardware Support

This feature supports IOM3 and above and applies only to the 7450 ESS and 7750 SR.

Any IOM3-supporting mode are subjected to the following check.

To launch cpe-ping on an Epipe, all of the following must be true:

1. All SAPs on the Epipe must be provisioned on slots that are mode-d compatible.

2. If bound to a PBB tunnel, all SAPs on the B-VPLS must be provisioned on slots that are mode-d compatible.
3. If the Epipe or the B-VPLS (in the case of PBB Epipe) uses SDP-bindings then the system configuration must be network-chassis-mode-d compatible.

3.2.24 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.

3.2.25 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.

3.2.26 VLL Diagnostics

3.2.27 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 and GRE SDP.

3.2.27.1 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:

- 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 7450 ESS, 7750 SR, and 7950 XRS routers.
- Use of the OAM control word as shown:

0 1 2 3

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| 0 0 0 1 | FmtID |   Reserved   |           Channel Type           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

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 PE node will only advertise the router alert label as the CC capability in the Label Mapping message. This method is supported by the 7450 ESS, 7750 SR and 7950 XRS routers.

- 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 7450 ESS, 7750 SR, and 7950 XRS routers.

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 below:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 0
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           0x0c           |           0x04           |   CC Types   |   CV Types   |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

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
- 0x02 MPLS Router Alert Label
- 0x04 MPLS inner label TTL = 1

If both PE nodes support more than one of the CC types, then the router 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 or GRE SDP and as such is not supported by the 7450 ESS, 7750 SR, and 7950 XRS routers.

0x02 LSP ping. This is used in VCCV ping application and applies to a VLL over an MPLS or a GRE SDP. This is supported by the 7450 ESS, 7750 SR, and 7950 XRS routers.


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 routers.
2. Reply via an IPv4/IPv6 UDP packet. This mode is supported by the routers.
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 routers.
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 routers.

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 router 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 router nodes. The VCCV ping feature can test connectivity of a VLL with any third party node which is compliant to RFC 5085.

Figure 25 VCCV-Ping Application

3.2.27.2 VCCV Ping in a Multi-Segment Pseudowire

Figure 26 shows an example of an application of VCCV ping over 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. The Switching PE (S-PE) is the node which performs pseudowire switching by cross-connecting two spoke SDPs.


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

- 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 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-vccv.

Note that the originator of the VCCV ping message does not need to be a T-PE node; it can be an S-PE node. The destination of the VCCV ping message can also be an S-PE node.

VCCV trace to trace the entire path of a pseudowire with a single command issued at the T-PE. This is equivalent to LSP trace and is an iterative process by which T-PE1 sends successive VCCV ping messages while incrementing the TTL value, starting from TTL=1. The procedure for each iteration is the same as above and each node in which the VC label TTL expires checks the FEC and replies with the FEC to the downstream S-PE or T-PE node. The process is terminated when the reply is from T-PE2 or when a timeout occurs.

Figure 26 VCCV Ping over a Multi-Segment Pseudowire



OSSG113

3.2.28 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-vccv, *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](#). 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 up to 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.

3.2.28.1 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.

3.2.28.2 Detailed VCCV-Trace Operation

[Figure 26](#) shows how 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 pseudo-wire 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.

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

3.2.28.3.1 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 router implementation will always make use of the user configuration for these parameters.

When in the trace mode operation, the T-PE will automatically learn the target FEC by probing one by one the hops of the MS-pseudowire path. Each S-PE node includes the FEC to the downstream node in the echo reply message in a similar way that LSP trace will have the probed node return the downstream interface and label stack in the echo reply message.

3.2.28.3.2 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.

3.2.28.3.3 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.

However, in trace mode, the node builds and sends the subsequent VCCV echo request with a incrementing TTL and the information (such as the downstream FEC 128) it received in the echo request to the next downstream pseudowire segment.

3.2.29 IGMP Snooping Diagnostics

3.2.30 MFIB Ping

The multicast forwarding information base (MFIB) ping OAM tool allows to easily verify inside a VPLS which SAPs would normally egress a certain multicast stream. The multicast stream is identified by a source unicast and destination multicast IP address, which are mandatory when issuing an MFIB ping command.

An MFIB ping packet will be sent through the data plane and goes out with the data plane format containing a configurable VC label TTL. 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 (SAP), it is identified by the OAM label below the VC label and passed to the management plane.

3.2.31 ATM Diagnostics

ATM Diagnostics applies to the 7450 ESS and 7750 SR only.

The ATM OAM ping allows operators to test VC-integrity and endpoint connectivity for existing PVCCs using OAM loopback capabilities.

If portId:vpi/vci PVCC does not exist, a PVCC is administratively disabled, or there is already a ping executing on this PVCC, then this command returns an error.

Because oam atm-ping is a dynamic operation, the configuration is not preserved. The number of oam atm-ping operations that can be performed simultaneously on a 7750 SR or 7450 ESS is configurable as part of the general OAM MIB configuration.

An operator can specify the following options when performing an oam atm-ping:

- **end-to-end** – this option allows sending oam atm-ping towards the connection endpoint in the line direction by using OAM end-to-end loopback cells
- **segment** – this option allows sending oam atm-ping towards the segment termination point in the line direction by using OAM segment loopback cells.

The result of ATM ping will show if the ping to a given location was successful. It also shows the round-trip time the ping took to complete (from the time the ping was injected in the ATM SAR device until the time the ping response was given to S/W by the ATM SAR device) and the average ping time for successful attempts up to the given ping response.

An oam atm ping in progress will time-out if a PVCC goes to the operational status down as result of a network failure, an administrative action, or if a PVCC gets deleted. Any subsequent ping attempts will fail until the VC's operational state changes to up.

To stop a ping in progress, an operator can enter "CTRL – C". This will stop any outstanding ping requests and will return ping result up to the point of interruption (a ping in progress during the above stop request will fail).

3.2.32 MPLS-TP On-Demand OAM Commands

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).

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

The 7450 ESS, 7750 SR, and 7950 XRS routers support VCCV Ping and VCCV Trace on single segment PWs and multi-segment PWs where every segment has static labels and a configured MPLS-TP PW Path ID. It also supports VCCV Ping and Trace on MS-PWs where a static MPLS-TP PW segment is switched to a dynamic T-LDP signaled segment.

Static MS-PW PWs are referred to with the sub-type static in the vccv-ping and vccv-trace command. 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.

Two ACH channel types are supported: the IPv4 ACH channel type, and the non-IP ACH channel type (0x0025). This is known as the non-ip associated channel. This is the default for type static. The Generic ACH Label (GAL) is not supported for PWs.

If the IPv4 associated channel is specified, 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 is always assumed to be the same application level control channel type for type static.

As with other PW types, the downstream mapping and detailed downstream mapping TLVs (DSMAP/DDMAP TLVs) are not supported on static MPLS-TP PWs.

The follow CLI command description shows the options that are only allowed if the type static option is configured. All other options are blocked.

```
vccv-ping static <sdp-id:vc-id> [target-fec-type pw-id-fec sender-src-address <ip-
address> remote-dst-address <ip-address> pw-id <value> pw-type <value>] [dest-
global-id <global-id> dest-node-id <node-id>] [assoc-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> [assoc-channel < ipv4 | non-ip>] [src-ip-address  
<ipv4-address>] [target-fec-type pw-id sender-src-address <ip-address> remote-  
dst-address <ip-address> pw-id <value> pw-type <value> ] [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>]
```

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 is 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 (i.e. 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.

3.2.33.1 VCCV Ping and VCCV Trace Between Static MPLS-TP and Dynamic PW Segments

The 7450 ESS, 7750 SR, and 7950 XRS routers support end to end VCCV Ping and VCCV trace between a segment with a static MPLS-TP PW and a dynamic T-LDP segment by allowing the user to specify a target FEC type for the VCCV echo request message that is different from the local segment FEC type. That is, it is possible to send a VCCV Ping / Trace echo request containing a static PW FEC in the target stack TLV at a T-PE where the local egress PW segment is signaled, or a VCCV Ping / Trace echo request containing a PW ID FEC (FEC128) in the target stack TLV at a T-PE where the egress PW segment is a static MPLS-TP PW.

Note that all signaled T-LDP segments and the static MPLS-TP segments along the path of the MS-PW must use a common associated channel type. Since only the IPv4 associated channel is supported in common between the two segments, this must be used. If a user selects a non-IP associated channel on the static MPLS-TP spoke-sdp, then vccv-ping and vccv-trace packets will be dropped by the S-PE.

The `target-fec-type` option of the `vccv-ping` and `vccv-trace` command is used to indicate that the remote FEC type is different from the local FEC type. For a `vccv-ping` initiated from a T-PE with a static PW segment with MPLS-TP parameters, attempting to ping a downstream FEC128 segment, then a `target-fec-type` of `pw-id` is configured with a static PW type. In this case, an `assoc-channel` type of `non-ip` is blocked, and vice-versa. Likewise the `reply-mode` must be set to `control-channel`. For a `vccv-ping` initiated from a T-PE with a FEC128 PW segment, attempting to ping a downstream static PW FEC segment, a `target-fec-type` of `static` is configured with a `pw-id` PW type, then a `control-channel` type of `non-ip` is blocked, and vice-versa. Likewise the `reply-mode` must also be set to `control-channel`.

When using VCCV Trace, where the first node to be probed is not the first-hop S-PE, the initial TTL must be set to `>1`. In this case, the `target-fec-type` refers to the FEC at the first S-PE that is probed.

The same rules apply to the `control-channel` type and `reply-mode` as for the `vccv-ping` case.

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

For `lsp-ping` and `lsp-trace` commands:

- 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 7450 ESS, 7750 SR, and 7950 XRS routers support the use of the G-ACh with non-IP encapsulation, IPv4 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]
```



```
[assoc-channel none | non-ip | ipv4] [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>]
  [assoc-channel none | non-ip | ipv4]
[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.

assoc-channel: If this is set to none, then IP encapsulation over an LSP is used with a destination address in the 127/8 range. If this is set to `ipv4`, then IPv4 encapsulation in a G-ACh 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 **config>saa>test**.

3.2.35 VxLAN Ping Supporting EVPN for VxLAN

EVPN is an IETF technology per RFC7432 that uses a new BGP address family and allows VPLS services to be operated as IP-VPNs, where the MAC addresses and the information to setup the flooding trees are distributed by BGP. The EVPN VxLAN connections, VxLAN Tunnel Endpoint (VTEP), uses a connection specific OAM Protocol for on demand connectivity verification. This connection specific OAM tool, VxLAN Ping, is described in the Layer 2 Services Guide, within the VxLAN Section.

3.2.36 Show Commands

The sample outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

3.2.37 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 shown below:

```
*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 Remote Address/Info	State Protocols	Tx Intvl Tx Pkts	Rx Intvl Rx Pkts	Multipl 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
=====

3.3 IP Performance Monitoring (IP PM)

The SR 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 \(TWAMP\)](#)
- [Two-Way Active Measurement Protocol Light \(TWAMP Light\)](#)

3.3.1 Two-Way Active Measurement Protocol (TWAMP)

Two-Way Active Measurement Protocol (TWAMP) provides a standards-based method for measuring the IP performance (packet loss, delay, and jitter) between two devices. TWAMP leverages 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, and 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”). The router acts as the “server”.

The control-client and server establish a TCP connection and exchange TWAMP-Control messages over this connection. When a server accepts the TCP control session from the control-client, it responds with a server greeting message. This greeting includes the various modes supported by the server. The modes are a bit mask. Each bit in the mask represents a functionality supported on the server. When the control-client wants to start testing, the client communicates the test parameters to the server, requesting any of the modes that the server supports. If the server agrees to conduct the described tests, the test begins as soon as the client sends a Start-Sessions or Start-N-Session 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. The exchange of test PDUs is referred to as TWAMP Test.

The TWAMP test PDU does not achieve symmetrical packet size in both directions unless the frame is padded by a minimum of 27 bytes. The session-sender is responsible for applying the desired padding. When an appropriate amount of padding is added, the session-reflector reduces the padding by the number of bytes needed to provide symmetry.

Server mode support includes:

- Individual Session Control (Mode Bit 4 : Value 16)
- Reflected Octets (Mode Bit 5 : Value 32)
- Symmetrical Size Test Packet (Mode Bit 6 : Value 64)

3.3.2 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 and synthetic loss 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 SR OS implementation supports the TWAMP Light model for gathering delay and loss statistics.

For TWAMP Light, the complete TWAMP model is replaced with a simple session-sender session-reflector.

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 reflector 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. Prior to release 13.0r4, if the configured TWAMP Light reflector UDP listening port was in use by another application on the system, a minor OAM message was presented indicating the UDP port was unavailable and that activation of the reflector is not allowed.

Notes: The TWAMP Light Reflector **udp-port** *udp-port-number* range configured as part of the **config>service|router>twamp-light create** command implements a restricted reserved UDP port range that must be adhere to range [64364..64373] prior to an upgrade or reboot. Configurations outside of this range will result in a failure of the TWAMP Light reflector or the prevention of the upgrade operation. If an In Service Software Upgrade (ISSU) function is invoked and the **udp-port** *udp-port-number* range is outside of the allowable range and the TWAMP Light Reflector is in a **no shutdown** state, the ISSU operation will not be allowed to proceed until, at a minimum, the TWAMP Light Reflector is **shutdown**. If the TWAMP Light Reflector is **shutdown**, the ISSU will be allowed to proceed, but the TWAMP Light Reflector will not be allowed to activate with a **no shutdown** until the range is brought in line the allowable range. A non-ISSU upgrade will be allowed to proceed regardless of the state (**shutdown** or **no shutdown**) of the TWAMP Light Reflector. The configuration will be allowed to load, but the TWAMP Light Reflector will remain inactive following the reload when the range is outside the allowable range. When the **udp-port** *udp-port-number* for a TWAMP Light Reflector is modified, all tests that were using the services of that reflector must update the **dest-udp-port** *udp-port-number* configuration parameter to match the new reflector listening port.

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. Configured prefixes can be modified without shutting down the reflector function. An inactivity timeout under the **config>oam-test>twamp>twamp-light** command 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 and beyond hardware.

Launching TWAMP Light test packets is under the control of the OAM Performance Monitoring (OAM-PM) architecture and as such adheres to those rules. This functionality is not available through interactive CLI or interactive SNMP, it is only available under the OAM-PM configuration construct. OAM-PM will report TWAMP Light delay and loss metrics. The OAM-PM architecture includes the assignment of a Test-ID. This protocol does not carry the 4-byte test ID in the packet. This is for local significance and uniformity with other protocols under the control of the OAM-PM architecture. The OAM-PM construct allows 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 UDP listening port on the reflector), the source UDP port and a number of other parameters that allow the operator to influence the packet handling. The source UDP port should only be configured when TWAMP distributed mode is being deployed. The probe interval and TWAMP Light packet padding size can be configured under the specific session. The pad size, the size of the all 0's pad, can configured to ensure that the TWAMP packet

is the same size in both directions. The session-sender role facilitated by the OAM-PM TWAMP Light testing will only set the multiplier bits in the Error Estimate field contained in the TWAMP test packet. The 8-bit multiplier field will be set to 00000001. The preceding 8 bits of the Error Estimate field comprised of S (1 bit - Time Sync), Z (1 bit MBZ) and Scale (6 bits) will all be set to 0.

TWAMP Test uses a single packet to gather both delay and loss metrics. This means there is special consideration over those approaches that utilize a specific tool per metric type.

In the TWAMP-Light case the interval parameter, which defines the probe spacing, is a common option applicable to all metrics collected under a single session. This requires the parameter to be removed from any test specific configurations, like the timing parameter associated with loss, specifically availability. Packet processing marks all fields in the PDU to report both delay and loss. The **record-stats** option can be used to refine which fields to process as part of the OAM-PM architecture. The default collection routine includes delay field processing only, **record-stats** delay. This is to ensure backward compatibility with previous releases that only supported the processing delay fields in the PDU. Enabling the processing of loss information requires the modification of the **record-stats** parameter. Adding loss to an active test requires the active test to be **shutdown**, modified and activate with the **no shutdown** command. It is critical to remember that the no shutdown action clears all previously allocated system memory for every test. Any results not written to flash or collected through SNMP are lost.

The **record-stats** setting do not change the configuration validation logic when a test is activated with the no shutdown command. Even if the loss metrics are not being processed and reported the configuration logic must ensure that the TWAMP test parameters are within the acceptable configuration limits, this includes default loss configuration statements. An operator has the ability to configure a TWAMP Light interval of 10s (10000ms) and record only delay statistics. The default **timing** parameter, used to compute and report availability and reliability, should allow for the activation of the test without a configuration violation. This requires the **frame-per-delta-t frames** default value of 1. An availability window cannot exceed 100s regardless of the **record-stats** setting. Computing the size of the availability window is a product of (**interval*frames-per-delta-t*consec-delta-t**).

The statistics display for the session with show all statistics that are being collected based on the **record-stats** configuration. If either of the metrics is not being recorded the statistics will display NONE for the excluded metrics.

Multiple tests sessions between peers are allowed. These test sessions are unique entities and may have different properties. Each test will generate TWAMP packets specific to their configuration.

TWAMP Light is supported on deployments that use IPv4 or IPv6 addressing, which may each have their own hardware requirements. All IP addressing must be unicast. IPv6 addresses can not be a reserved or a link local address. Multiple test sessions may be configured between the same source and destination IP endpoints. The tuple Source IP, Destination IP, Source UDP, and Destination UDP provide a unique index for each test point.

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

See the OAM-PM section of this guide for more information about the integration of TWAMP Light and the OAM-PM architecture, including hardware dependencies.

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.

The following examples 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 64364 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 64364
  destination 10.1.1.1
  fc "l2"
(default) no forwarding
  profile in
  router 500
  source 10.2.1.1
(default) ttl 255
  twamp-light test-id 500 create
(default) interval 1000
  loss
(default) flr-threshold 50
(default) timing frames-per-delta-t 1 consec-delta-t 10 chli-
threshold 5
  exit
  pad-size 27
(default) record-stats delay
  no test-duration
  no shutdown
  exit
exit
-----
```

3.4 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 up to 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 SR and ESS platforms.

This section of the guide will provide configuration example for each of the functions. It will also provide the various OAM command line options and show commands to operate the network. The individual service guides will provide the complete CLI configuration and description of the commands in order to build the necessary constructs and management points.

Acronym	Callout	Supported platform
1DM	One way Delay Measurement (Y.1731)	All
AIS	Alarm Indication Signal	All
CCM	Continuity check message	All
CFM	Connectivity fault management	All
CSF	Client Signal Fail (Receive)	All
DMM	Delay Measurement Message (Y.1731)	All
DMR	Delay Measurement Reply (Y.1731)	All
LBM	Loopback message	All
LBR	Loopback reply	All
LMM	(Frame) Loss Measurement Message	Platform specific

Acronym	Callout (Continued)	Supported platform
LMR	(Frame) Loss Measurement Response	Platform specific
LTM	Linktrace message	All
LTR	Linktrace reply	All
ME	Maintenance entity	All
MA	Maintenance association	All
MA-ID	Maintenance association identifier	All
MD	Maintenance domain	All
MEP	Maintenance association end point	All
MEP-ID	Maintenance association end point identifier	All
MHF	MIP half function	All
MIP	Maintenance domain intermediate point	All
OpCode	Operational Code	All
RDI	Remote Defect Indication	All
TST	Ethernet Test (Y.1731)	All
SLM	Synthetic Loss Message	All
SLR	Synthetic Loss Reply (Y.1731)	All

This section describes the following topics:

- [ETH-CFM Building Blocks](#)
- [Loopback](#)
- [Loopback Multicast](#)
- [Linktrace](#)
- [Continuity Check \(CC\)](#)
- [CC Remote Peer Auto-Discovery](#)
- [CCM Grace Period](#)
- [CCM Hold Timers](#)
- [Alarm Indication Signal \(ETH-AIS Y.1731\)](#)
- [Client Signal Fail \(ETH-CSF Y.1731\)](#)
- [Test \(ETH-TST Y.1731\)](#)
- [One-Way Delay Measurement \(ETH-1DM Y.1731\)](#)
- [Two-Way Delay Measurement \(ETH-DMM Y.1731\)](#)
- [Synthetic Loss Measurement \(ETH-SLM Y.1731\)](#)
- [Frame Loss Measurement \(ETH-LMM Y.1731\)](#)

3.4.1 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 SR 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 up to 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, up to 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) from 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 SR OS 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.

In the following example, a Y.1731 domain context and 802.1ag context are configured. The Y.1731 context can be identified by the **none** setting for the domain format.

```
configure eth-cfm domain 3 format none level 3
configure eth-cfm domain 4 format string name IEEE-Domain level 4
```

The chassis does not support a domain format of **none** for the 802.1ag contexts. The domain index, the first numerical value, is not related to the level, even though in this example they do match.

The following example illustrates the creation of the association within the domain context. The association links the construct to the service using the value of the bridge-identifier. The value specified for the bridge-identifier is equivalent to the numerical value used to create the service.

```
config>eth-cfm# info
-----
      domain 3 format none level 3
        association 1 format icc-based name "123456789abcd"
          bridge-identifier 100
        exit
      exit
        association 2 format string name "Y1731ContextIEEEFormat"
          bridge-identifier 300
        exit
      exit
    exit
  domain 4 name "IEEE-Domain" level 4
    association 1 format string name "UpTo45CharactersForIEEEString"
      bridge-identifier 100
    exit
```

```
        ccm-interval 1
    exit
exit
-----
```

This example show how to format the association within the domain to match the domain format, Y.1731 (domain 3/association 1) or 802.1ag (domain 4/association 1), and how the 802.1ag association format can be configured within a Y.1731 domain (domain 3/association 2). The mixed configuration represented by domain 3 association 2 may be of value in mixed Y.1731 and 802.1ag environments.

The CCM-interval is also specified within the association and has a default of 10 seconds unless specifically configured with another value. When the association is created and the MEP is a facility MEP the bridge-identifier is not to be included in the configuration since the facility MEP is not bound to a service. Facility MEPs are described in the *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide*.

Maintenance Endpoints/MEG Endpoints (MEP) are the workhorses of ETH-CFM. A MEP is the unique identification within the association (1-8191). 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. A vMEP is a service level MEP configuration that installs ingress (down MEP-like) extraction on the supported ETH-CFM termination points within a VPLS configuration.

Maintenance Intermediate Points/MEG Intermediate Points (MIPs) are management entities between the terminating MEPs along the service path. MIPs 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.

MIP creation is the result of the **mhf-creation** mode and interaction with related MEPs, and with the direction of the MEP. Two different authorities can be used to determine the MIPs that should be considered and instantiated. The domain and association or the default-domain hierarchies match the configured bridge identifier and VLAN to the service ID and any configured primary VLAN. When a primary VLAN MIP is not configured, the VLAN is either ignored or configured as **none**. The

mhf-creation command is described in the *7450 ESS, 7750 SR, and 7950 XRS Services Overview Guide*.

The domain and association MIP creation function triggers a search for all ETH-CFM domain association bridge identifier matches to the service ID. A MIP candidate is then be evaluated using the **mhf-creation** mode and the rules that govern the algorithm. The domain association **mhf-creation** modes and their uses are listed below:

- none: a MIP is not a candidate for creation using this domain association bridge identifier. This is the default **mhf-creation** mode for every bridge identifier under this hierarchy.
- explicit: a MIP is a candidate for creation using this domain association bridge identifier only if a lower-level MEP exists.
- default: a MIP is a candidate for creation using this domain association bridge identifier regardless of the existence of a lower-level MEP. If a lower-level MEP is present, this creation mode behaves in the same manner as explicit creation mode.
- static: a MIP is a candidate for creation using the domain association bridge identifier at the level of the domain. This creation mode is specific to MIPs with the **primary-vlan-enabled** parameter configured. Different VLANs maintain their own level hierarchies. Primary VLAN creation under this context requires static mode.

For all modes except static mode, only a single MIP can be created. All candidates are collected and the lowest-level valid MIP is created. In static mode, all valid MIPs will be created for the bridge identifier VLAN pair. A MIP is considered invalid if the level of the MIP is equal to or below a downward-facing MEP, or below the level of an upward-facing MEP and the MIP shares the same service component as the Up MEP.

Not all creation modes require the **mip** creation statement within the service. The explicit and default **mhf-creation** modes may instantiate a MIP without the **mip** creation statement under the service if a lower-level MEP exists for the domain association bridge identifier. If a lower-level MEP does not exist, the default and static **mhf-creation** modes require the **mip** creation statement on the service connection.

MEPs require the domain and association configurations to ensure that all ETH-CFM PDUs can be supported. MIPs have restricted ETH-CFM PDU support: ETH-LB and ETH-LT. These two protocols do not require the configuration of a domain and association. MIPs may be created outside of the association context using the default-domain table.

The **default-domain** table is an object table populated with values that are used for MIP creation. The table is indexed by the bridge identifier and VLAN. An index entry is automatically added when the **mip** creation statement is added under a SAP or SDP binding. When an index entry is added, the bridge identifier is set to the service ID and the VLAN is set to the **primary-vlan-enable** *vlan-id*. If the MIP does not use primary VLAN functionality, the VLAN will be configured as **none**. When the entry has been added to the default-domain table, the default values can be configured. The default-domain table defers to the system-wide, read-only values.

Because there are two different locations able to process the MIP creation logic, a per-bridge identifier VLAN authority must be determined. The authority is a component, table, or configuration that is responsible for executing the MIP creation algorithm. In general, any domain association bridge identifier that could be used to create a specific MIP is authoritative. Other configurations influence the authority, such as the type of MIP (primary VLAN or non-primary VLAN), the different **mhf-creation** modes, the interaction of those modes with MEPs, and the direction of the MEP.

The following rules provide some high-level guidelines to determine the authority.

- Rule 1: The original model predating the **default-domain** is always applied first. If a MIP is created using the original model, the new model will not be applied. The original model includes complex Up MEP MIP creation rules. If an Up MEP exists on a service connection, any service connection other than the one with the active Up MEP will attempt to create the lowest higher-level MIP using the domain association bridge identifier table. If a higher-level MIP cannot be created, and no higher-level association exists, the default-domain table is consulted.
- Rule 2: A **mip** creation statement is required under the service connection in order to use the default-domain table. This is different from the domain association table. The domain association table does not require the **mip** creation statement when the **mhf-creation** mode is configured as either explicit or default and a lower-level MEP is present.
- Rule 3: If no domain association bridge identifier matches the service ID, the default-domain table is consulted.
- Rule 4: If a domain association bridge identifier matches a service ID for the sole purpose of MEP creation, and no higher or lower domain association with the same bridge identifier exists, the default-domain table is consulted.
 - Rule 4a: Any domain association bridge identifier matching a service ID with a configured VLAN and a static **mhf-creation** mode is authoritative for all matching service IDs and MIPs with **primary-vlan-enable** configured with the same VLAN.

- Rule 4b: Any domain association bridge identifier attempting to create a MIP with **primary-vlan-enable** configured is considered non-authoritative if the **mhf-creation** mode is anything other than static.

When the authority for MIP creation is determined, the MIP attributes are derived from that creation table. The default-domain table defers to the read-only, system-wide MIP values and inherits those defaults. Some of the objects under the default-domain hierarchy must be configured using the same statement to avoid transient and unexpected MIP creation while the configuration is being completed. To this end, the **mhf-creation** mode and level have been combined in the same configuration statement.

The standard **mhf-creation** modes (none, default, explicit) are configurable as part of the default-domain table. Static mode can only be configured under the domain association bridge identifier. This is because default-domain table indexing precludes multiple MIPs at different levels.

MIP creation requires configuration. The default values in both the domain association and the default-domain table prevent MIP instantiation.

The **show eth-cfm mip-instantiation** command can be used to check the authority for each MIP.

There are two locations in the configuration where ETH-CFM is defined. The domains, associations (including links to the service ID), MIP creation method, common ETH-CFM functions, and remote MEPs are defined under the top-level **eth-cfm** command. If Y.1731 functions are required, the context under which the MEPs are configured should follow the Y.1731 specific formats (domain format of none). The Y.1731 functions (**ais-enable**, **eth-test-enable**, and **csf-enable**) are rejected under a non-Y.1731 MEP. The Y.1731 specific on-demand tests are blocked from the global CLI. The Y.1731 PM tests are not blocked when executed against a non-Y.1731 domain format when launched from OAM-PM and SAA. When these parameters have been entered, the MEP, and possibly the MIP, can be defined within the service under the SAP or SDP binding.

[Table 6](#) is a general table that indicates ETH-CFM support for the different services and SAP or SDP binding. 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**

Service	Ethernet Connection	Down MEP	Up MEP	MIP	Virtual MEP
Epipe	—	—	—	—	No
	SAP	Yes	Yes	Yes	—
	Spoke-SDP	Yes	Yes	Yes	—
VPLS	—	—	—	—	Yes
	SAP	Yes	Yes	Yes	—
	Spoke-SDP	Yes	Yes	Yes	—
	Mesh-SDP	Yes	Yes	Yes	—
B-VPLS	—	—	—	—	Yes
	SAP	Yes	Yes	Yes	—
	Spoke-SDP	Yes	Yes	Yes	—
	Mesh-SDP	Yes	Yes	Yes	—
I-VPLS	—	—	—	—	No
	SAP	Yes	Yes	Yes	—
	Spoke-SDP	Yes	Yes	Yes	—
M-VPLS	—	—	—	—	No
	SAP	Yes	Yes	Yes	—
	Spoke-SDP	Yes	Yes	Yes	—
	Mesh-SDP	Yes	Yes	Yes	—
PBB EPIPE	—	—	—	—	No
	SAP	Yes	Yes	Yes	—
	Spoke-SDP	Yes	Yes	Yes	—
IPIPE	—	—	—	—	No
	SAP	Yes	No	No	—
	Ethernet-Tunnel SAP	Yes	No	No	—

Table 6 ETH-CFM Support Matrix (Continued)

Service	Ethernet Connection	Down MEP	Up MEP	MIP	Virtual MEP
IES	—	—	—	—	No
	SAP	Yes	No	No	—
	Spoke-SDP (Interface)	Yes	No	No	—
	Subscriber Group-int SAP	Yes	No	No	—
VPRN	—	—	—	—	No
	SAP	Yes	No	No	—
	Spoke-SDP (Interface)	Yes	No	No	—
	Subscriber Group-int SAP	Yes	No	No	—



Figure 27 MEP and MIP

Figure 28 illustrates the usage of an EPIPE on two different nodes that are connected using ether SAP 1/1/2:100.31. The SAP 1/1/10:100.31 is an access port that is not used to connect the two nodes.

Figure 28 MEP Creation



```

NODE1
config>eth-cfm# info
-----
    domain 3 format none level 3
      association 1 format icc-based name "03-0000000101"
      bridge-identifier 100
      exit
    exit
  exit
  domain 4 format none level 4
    association 1 format icc-based name "04-0000000102"
    bridge-identifier 100
    exit
  exit
exit

config>service>epipe# info
-----
    sap 1/1/2:100.31 create
      eth-cfm
        mep 111 domain 3 association 1 direction down
        mac-address d0:0d:1e:00:01:11
        no shutdown
      exit
    exit
  exit
  sap 1/1/10:100.31 create
    eth-cfm
      mep 101 domain 4 association 1 direction up
      mac-address d0:0d:1e:00:01:01
      no shutdown
    exit
  exit

```

```

        exit
        no shutdown
    -----

NODE 2
eth-cfm# info
-----
    domain 3 format none level 3
        association 1 format icc-based name "03-0000000101"
            bridge-identifier 100
            exit
        exit
    exit
    domain 4 format none level 4
        association 1 format icc-based name "04-0000000102"
            bridge-identifier 100
            exit
        exit
    exit
    exit
-----

config>service>epipe# info
-----
    sap 1/1/2:100.31 create
        eth-cfm
            mep 112 domain 3 association 1 direction down
            mac-address d0:0d:1e:00:01:12
            no shutdown
            exit
        exit
    exit
    sap 1/1/10:100.31 create
        eth-cfm
            mep 102 domain 4 association 1 direction up
            mac-address d0:0d:1e:00:01:02
            no shutdown
            exit
        exit
    exit
    exit
    no shutdown
    -----


```

Examining the configuration from NODE1, MEP 101 is configured with a direction of UP causing all ETH-CFM traffic originating from this MEP to generate into the switch fabric and out the mate SAP 1/1/2:100.31. MEP 111 uses the default direction of DOWN causing all ETH-CFM traffic that is generated from this MEP to send away from the fabric and only egress the SAP on which it is configured, SAP 1/1/2:100.31.

Further examination of the domain constructs reveal that the configuration properly uses domain nesting rules. In this case, the Level 3 domain is completely contained in a Level 4 domain.

[Figure 29](#) illustrates the creation of and explicit MIP.

Figure 29 MIP Creation Example (NODE1)



OSSG549

```

NODE1
config>eth-cfm# info
-----
    domain 3 format none level 3
        association 1 format icc-based name "03-0000000101"
            bridge-identifier 100
            exit
        exit
    exit
    domain 4 format none level 4
        association 1 format icc-based name "04-0000000102"
            bridge-identifier 100
            exit
        exit
    association 2 format icc-based name "04-MIP0000102"
        bridge-identifier 100
        mhf-creation explicit
        exit
    exit
    exit

config>service>epipe# info
-----
    sap 1/1/2:100.31 create
        eth-cfm
            mep 111 domain 3 association 1 direction down
                mac-address d0:0d:1e:00:01:11
                no shutdown
            exit
        exit
    exit
    sap 1/1/10:100.31 create
        eth-cfm
            mep 101 domain 4 association 1 direction up
                mac-address d0:0d:1e:00:01:01
                no shutdown
        exit
    exit

```

```


        exit
    exit
    exit
    no shutdown
-----
NODE 2
eth-cfm# info
-----
    domain 3 format none level 3
        association 1 format icc-based name "03-0000000101"
            bridge-identifier 100
        exit
    exit
    exit
    domain 4 format none level 4
        association 1 format icc-based name "04-0000000102"
            bridge-identifier 100
        exit
    exit
    association 2 format icc-based name "04-MIP0000102"
        bridge-identifier 100
        mhf-creation explicit
    exit
    exit
    exit
-----
config>service>epipe# info
-----
    sap 1/1/2:100.31 create
        eth-cfm
            mep 112 domain 3 association 1 direction down
            mac-address d0:0d:1e:00:01:12
            no shutdown
        exit
    exit
    exit
    sap 1/1/10:100.31 create
        eth-cfm
            mep 102 domain 4 association 1 direction up
            mac-address d0:0d:1e:00:01:02
            no shutdown
        exit
    exit
    exit
    no shutdown
-----

```

An addition of association 2 under domain four includes the **mhf-creation explicit** statement. This means that when the level 3 MEP is assigned to the SAP 1/1/2:100.31 using the definition in domain 3 association 1, creating the higher level MIP on the same SAP. Since a MIP does not have directionality “Both” sides are active. The service configuration and MEP configuration within the service did not change.

Figure 30 illustrates a simpler method that does not require the creation of the lower level MEP. The operator simply defines the association parameters and uses the **mhf-creation default** setting, then places the MIP on the SAP of their choice.

Figure 30 MIP Creation Default



OSSG550

NODE1:

```
config>eth-cfm# info
-----
domain 4 format none level 4
  association 1 format icc-based name "04-0000000102"
    bridge-identifier 100
    exit
  exit
  association 2 format icc-based name "04-MIP0000102"
    bridge-identifier 100
    mhf-creation default
    exit
  exit
exit

config>service>epipe# info
-----
sap 1/1/2:100.31 create
  eth-cfm
    mip mac d0:0d:1e:01:01:01
  exit
exit
sap 1/1/10:100.31 create
  eth-cfm
    mep 101 domain 4 association 1 direction up
    mac-address d0:0d:1e:00:01:01
    no shutdown
  exit
exit
exit
```

```

no shutdown
-----

NODE2:

config>eth-cfm# info
-----
    domain 4 format none level 4
      association 1 format icc-based name "04-0000000102"
        bridge-identifier 100
        exit
      exit
      association 2 format icc-based name "04-MIP0000102"
        bridge-identifier 100
        mhf-creation default
        exit
      exit
    exit
  exit
-----

config>service>epipe# info
-----
    sap 1/1/2:100.31 create
      eth-cfm
        mip mac d0:0d:1e:01:01:02
        exit
      exit
    sap 1/1/10:100.31 create
      eth-cfm
        mep 102 domain 4 association 1 direction up
        mac-address d0:0d:1e:00:01:02
        no shutdown
        exit
      exit
    exit
  exit
no shutdown
-----

```

Figure 31 shows the detailed IEEE representation of MEPs, MIPs, levels and associations, using the standards defined icons.

SAPs support a comprehensive set of rules including wild cards to map packets to services. For example, a SAP mapping packets to a service with a port encapsulation of QinQ may choose to only look at the outer VLAN and wildcard the inner VLAN. SAP 1/1/1:100.* would map all packets arriving on port 1/1/1 with an outer VLAN 100 and any inner VLAN to the service the SAP belongs to. These powerful abstractions will extract inbound ETH-CFM PDUs only when there is an exact match to the SAP construct. In the case of the example when then an ETH-CFM PDU arrives on port

1/1/1 with a single VLAN with a value of 100 followed immediately with e-type (0x8902 ETH-CFM). Furthermore, the generation of the ETH-CFM PDUs that egress this specific SAP will be sent with only a single tag of 100. If the operator needs to extract ETH-CFM PDUs or generate ETH-CFM PDUs on wildcard SAPs Primary VLAN will be required.

Table 7 shows how packets that would normally bypass the ETH-CFM extraction would be extracted when Primary VLAN is configured. This assumes that the processing rules for MEPs and MIPs is met, E-type 0x8902, Levels and OpCodes.

Table 7 Extraction Comparison with Primary VLAN

Port Encapsulation	E-type	Ingress Tag(s)	Ingress SAP	No Primary VLAN ETH-CFM Extraction		With Primary VLAN (10) ETH-CFM Extraction	
				MEP	MIP	MEP	MIP
—	—	—	—	MEP	MIP	MEP	MIP
Dot1q	0x8902	10	x/y/z:*	No	No	Yes	Yes
Dot1q	0x8902	10.10	x/y/z:10	No	No	Yes	Yes
QinQ	0x8902	10.10	x/y/z:10.*	No	No	Yes	Yes
QinQ (Default Behavior)	0x8902	10.10	x/y/z:10.0	No	No	Yes	Yes
Null	0x8902	10	x/y/z	No	No	Yes	Yes

The mapping of the service data remains unchanged. The Primary VLAN function allows for one additional VLAN offset beyond the SAP configuration, up to a maximum of two VLANs in the frame. If a fully qualified SAP specifies two VLANs (SAP 1/1/1:10.10) and a primary VLAN of 12 is configured for the MEP there will be no extraction of ETH-CFM for packets arriving tagged 10.10.12. That exceeds the maximum of two tags.

The mapping of service data based on SAPs has not changed. ETH-CFM MP's functionality remains SAP specific. In instances where a service includes a specific SAP with a specified VLAN (1/1/1:50) and a wildcard SAP on the same port (1/1/1:*) it is important to understand how the ETH-CFM packets are handled. Any ETH-CFM packet with e-type 0x8902 arriving with a single tag or 50 would be mapped to a classic MEP configured under SAP 1/1/1:50. Any packet arriving with an outer VLAN of 50 and second VLAN of 10 would be extracted by the 1/1/1:50 SAP and would

require a Primary VLAN enabled MEP with a value of 10, assuming the operator would like to extract the ETH-CFM PDU of course. An inbound packet on 1/1/1 with an outer VLAN tag of 10 would be mapped to the SAP 1/1/1:*. If ETH-CFM extraction is required under SAP 1/1/1:* a Primary VLAN enabled MEP with a value of 10 would be required.

Obviously, the packet that is generated from a MEP or MIP with Primary VLAN enabled will include that VLAN. The SAP will encapsulate the Primary VLAN using the SAP encapsulation.


Primary VLAN support includes UP MEPs, DOWN MEPs and MIPs on Ethernet SAPs, including LAG for ePipe and VPLS services. There is no support for Primary VLAN configuration for vMEPs or MEPs on SDP binding. Classic MEPs, those without a primary VLAN enabled, and Primary VLAN enabled MEPs can co-exist under the same SAP. Classic MIPs and Primary VLAN enabled MIPs may also coexist. The enforcement of a single classic MIP per SAP continues to be enforced. However, the operator may configure multiple Primary VLAN enabled MIPs on the same SAP. MIPs in the Primary VLAN space must include the mhf-creation static under the association and must also include the specific VLAN on the MIP creation statement under the SAP. The **no** version of the **mip** command must include the entire statement including the VLAN information.

The eight MD Levels (0-7) are specific to context in which the Management Point (MP) is configured. This means the classic MPs have a discrete set of the levels from the Primary VLAN enabled space. Each Primary VLAN space has its own eight Level MD space for the specified Primary VLAN. Consideration must be given before allowing overlapping levels between customers and operators should the operator be provision a customer facing MP, like a MIP on a UNI. CPU Protection extensions for ETH-CFM are VLAN unaware and based on MD Level and the OpCode. Any configured rates will be applied to the Level and OpCode as a group.

There are two configuration steps to enable Primary VLAN. Under the bridging instance, contained within the association context (cfg>eth-cfm>domain>assoc>bridge) the VLAN information must be configured. Until this is enabled using the *primary-vlan-enable* option as part of the MEP creation step or the MIP statement (cfg>service>...>sap>eth-cfm>) the VLAN specified under the bridging instance remains inactive. This is to ensure backward interoperability.

Primary VLAN functions require a minimum of IOM3/IMM. There is no support for vpls-sap-templates. Sub second CCM intervals are not supported for Primary VLAN MEPs.

Figure 31 MEP, MIP and MD Levels




An operator may see the following INFO message (during configuration reload), or MINOR (error) message (during configuration creation) when upgrading to 11.0r4 or later if two MEPs are in a previously undetected conflicting configuration. The messaging is an indication that a MEP, the one stated in the message using format (domain <md-index> / association <ma-index> / mep <mep-id>), is already configured and has allocated that context. During a reload (INFO) a MEP that encounters this condition will be created but its state machine will be disabled. If the MINOR error occurs during a configuration creation this MEP will fail the creation step. The indicated MEP will need to be correctly re-configured.

```
INFO: ETH_CFM #1341 Unsupported MA ccm-interval for this MEP - MEP 1/112/
21 conflicts with sub-second config on this MA
MINOR: ETH_CFM #1341 Unsupported MA ccm-interval for this MEP - MEP 1/112/
21 conflicts with sub-second config on this MA
```

3.4.2 Loopback

A loopback message is generated by an MEP to its peer MEP or a MIP (Figure 32). The functions are similar to an IP ping to verify Ethernet connectivity between the nodes.

Figure 32 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. The ETH-LB PDU format does not allow a MEP to have more than a single active ETH-LB session.
- MIP — Supports responding to loopback messages with loopback reply messages when loopback messages are targeted to self
- SenderID TLV may optionally be configured to carry the ChassisID. When configured, this information will be included in LBM messages.
 - Only the ChassisID portion of the TLV will be included
 - The Management Domain and Management Address fields are not supported on transmission
 - As per the specification, the LBR function copies and returns any TLVs received in the LBM message. This means that the LBR message will include the original SenderID TLV.
 - Supported for both service (id-permission) and facility MEPs (facility-id-permission)
 - Supported for both MEP and MIP
- Displays the loopback test results on the originating MEP

The ETH-LBM (loopback) function includes parameters for sub second intervals, timeouts, and new padding parameters. The CLI display output has been enhanced to provide more information and a new format.

When an ETH-LBM command is issued using a sub second interval (100ms), the output success will be represented with a “!” character, and a failure will be represented with a “.” The updating of the display will wait for the completion of the previous request before producing the next result. However, the packets will maintain the transmission spacing based on the interval option specified in the command.

```
oam eth-cfm loopback 00:00:00:00:00:30 mep 28 domain 14 association 2 interval 1
send-count 100 timeout 1
Eth-Cfm Loopback Test Initiated: Mac-Address: 00:00:00:00:00:30, out service: 5

!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

Sent 100 packets, received 100 packets [0 out-of-order, 0 Bad Msdu]
Packet loss 1.00%
```

When the interval is one seconds or higher, the output will provide detailed information that includes the number of bytes (from the LBR), the source MEP ID (format md-index/ma-index/mepid), and the sequence number as it relates to this test and the result.

```
oam eth-cfm loopback 00:00:00:00:00:30 mep 28 domain 14 association 2 interval 10
send-count 10 timeout 1
Eth-Cfm Loopback Test Initiated: Mac-Address: 00:00:00:00:00:30, out service: 5

56 bytes from 14/2/28; lb_seq=1 passed
56 bytes from 14/2/28; lb_seq=2 passed
56 bytes from 14/2/28; lb_seq=3 passed
56 bytes from 14/2/28; lb_seq=4 passed
56 bytes from 14/2/28; lb_seq=5 passed
56 bytes from 14/2/28; lb_seq=6 passed
56 bytes from 14/2/28; lb_seq=7 passed
56 bytes from 14/2/28; lb_seq=8 passed
56 bytes from 14/2/28; lb_seq=9 passed
56 bytes from 14/2/28; lb_seq=10 passed

Sent 10 packets, received 10 packets [0 out-of-order, 0 Bad Msdu]
Packet loss 0.00%
```

Since ETH-LB does not support standard timestamps, no indication of delay is produced as these times are not representative of network delay.

By default, if no interval is included in the command, the default is back to back LBM transmissions. The maximum count for such a test is 5.

3.4.3 Loopback Multicast

Multicast loopback also support the new intervals. However, the operator must be careful when using this approach. Every MEP in the association will respond to this request. This means an exponential impact on system resources for large scale tests. If the multicast option is used and there with an interval of 1 (100ms) and there are 50 MEPs in the association, this will result in a 50 times increase in the receive rate (500pps) compared to a unicast approach. Multicast displays will not be updated until the test is completed. There is no packet loss percentage calculated for multicast loopback commands.

This on demand operation tool is used to quickly check the reachability of all MEPs within an Association. A multicast address can be coded as the destination of an **oam eth-cm loopback** command. The specific class 1 multicast MAC address or the keyword “multicast” can be used as the destination for the loopback command. The class 1 ETH-CFM multicast address is in the format 01:80:C2:00:00:3x (where x = 0 - 7 and is the number of the domain level for the source MEP). When the “multicast” option is used, the class 1 multicast destination is built according to the local MEP level initiating the test.

Remote MEPs that receive this message, configured at the equivalent level, will terminate and process the multicast loopback message responding with the appropriate unicast loopback response (ETH-LBR). Regardless of whether a multicast or unicast ETH-LBM is used, there is no provision in the standard LBR PDU to carry the MEP-ID of the responder. This means only the remote MEP MAC Address will be reported and subsequently displayed. MIPs will not respond to the multicast ETH-LBM. It is important to understand that although MIPs do not respond they perform the basic level and opcode check to determine whether they need to decode the packet. MIPs along the applicable path over which the LBM is sent that match the level and opcode will decode the packet, not respond and forward along the path.

MEP loopback stats are not updated as a result of this test being run. That means the received, out-of-order and bad-msdu counts are not affected by multicast loopback tests. The multicast loopback command is meant to provide immediate connectivity troubleshooting feedback for remote MEP reachability only.

```
oam eth-cfm loopback multicast mep 28 domain 14 association 2 interval 1 send-  
count 100  
Eth-Cfm Loopback Test Initiated: Mac-Address: multicast, out service: 5
```

MAC Address	Receive Order												
00:00:00:00:00:30	1	2	3	4	5	6	7	8	9	10	11	12	13
14	15	16	17	18	19	20	21	22	23	24	25	26	27
28	29	30	31	32	33	34	35	36	37	38	39	40	41
42	43	44	45	46	47	48	49	50	51	52	53	54	55
56	57	58	59	60	61	62	63	64	65	66	67	68	69


```
4 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80
81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97
98 99 100
00:00:00:00:00:32 1 2 3 4 5 6 7 8 9 10 11 12 13
14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30
31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47
48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 6
4 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80
81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97
98 99 100

Sent 100 multicast packets, received 200 packets
```

3.4.4 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 33](#)). Linktrace 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.


Figure 33 CFM Linktrace

Fig_13

The IEEE and ITU-T handle the linktrace reply slightly differently. An IEEE 802.1ag configured MEP requires the relay action field to be a valid non-zero integer. The ITU-T ignores the relay action field and will set the value to zero when responding to the LTM. In mixed 802.ag and Y.1731 environments the operator may choose to configure a Y.1731 context with an IEEE domain format.

The following linktrace related functions are supported:

- MEP — Supports generating linktrace messages and responding with linktrace reply messages. The ETH-LT PDU format does not allow a MEP to have more than a single active ETH-LT session.
- 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
- SenderID TLV may optionally be configured to carry the ChassisID. When configured, this information will be included in LTM and LTR messages.
 - Only the ChassisID portion of the TLV will be included
 - The Management Domain and Management Address fields are not supported on transmission
 - THE LBM message will include the SenderID TLV that is configured on the launch point. The LBR message will include the SenderID TLV information from the reflector (MIP or MEP) if it is supported.
 - Supported for both service (id-permission) and facility MEPs (facility-id-permission).
 - Supported for both MEP and MIP

The following output includes the SenderID TLV contents if it is included in the LBR.

```
oam eth-cfm linktrace 00:00:00:00:00:30 mep 28 domain 14 association 2
Index Ingress Mac          Egress Mac          Relay    Action
-----
1      00:00:00:00:00:00    00:00:00:00:00:30  n/a      terminate
SenderId TLV: ChassisId (local)
              access-012-west
-----
No more responses received in the last 6 seconds.
```

3.4.5 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 34 CFM Continuity Check






Figure 35 CFM CC Failure Scenario

An MEP may be configured to generate ETH-CC packet using a unicast destination Layer 2 MAC address. This may help reduce the overhead in some operational models where Down MEPs per peer are not available. For example, mapping an I-VPLS to a PBB core where a hub is responsible for multiple spokes is one of the applicable models. When ETH-CFM packets are generated from an I-context toward a remote I-context, the packets will traverse the B-VPLS context. Since many B-contexts are multipoint, any broadcast, unknown or multicast packet is flooded to all appropriate nodes in the B-context. When ETH-CC multicast packets are generated, all the I-VPLS contexts in the association must be configured with all the appropriate remote MEPids. If direct spoke to spoke connectivity is not part of the validation requirement, the operational complexity can be reduced by configuring unicast DA addressing on the “spokes” and continuing to use multicast CCM from the “hub”. When the unicast MAC is learned in the forwarding DB, traffic will be scoped to a single node.

Figure 36 Unicast CCM in Hub & Spoke Environments



Defect condition, reception, and processing will remain unchanged for both hub and spokes. When an ETH-CC defect condition is raised on the hub or spoke, the appropriate defect condition will be set and distributed throughout the association from the multicasting MEP. For example, should a spoke raise a defect condition or timeout, the hub will set the RDI bit in the multicast ETH-CC packet which is received on all spokes. Any local hub MEP defect condition will continue to be propagated in the multicast ETH-CC packet. Defect conditions will be cleared as per normal behavior.

The forwarding plane must be considered before deploying this type of ETH-CC model. A unicast packet will be handled as unknown when the destination MAC does not exist in local forwarding table. If a unicast ETH-CC packet is flooded in a multipoint context, it will reach all the appropriate I-contexts. This will cause the spoke MEPs to raise the “DefErrorCCM” condition because an ETH-CC packet was received from a MEP that has not been configured as part of the receiving MEPs database.

The remote unicast MAC address must be configured and is not automatically learned. A MEP cannot send both unicast and multicast ETH-CC packets. Unicast ETH-CC is only applicable to a local association with a single configured remote peer. There is no validation of MAC addresses for ETH-CC packets. The configured unicast destination MAC address of the peer MEP only replaces the multicast class 1 destination MAC address with a unicast destination.

Unicast CCM is not supported on any MEPs that are configured with sub second CCM-intervals.

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: 10ms, 100ms, 1s, 10s 60s, 600s. Default: 10s. Sub-second, or fast CC requires a 7950 XRS, 7750 SR-7/SR-12/SR-12e with a minimum SF/CPM-3, 7750 ESS-7/ESS-12 with a minimum SF/CPM-3, 7750 SRa, or 7750 SRe, and only a limited number supported on SF/CPM-1 and SF/CPM-2. When configuring MEPs with sub-second CCM intervals, bandwidth consumption must be taken into consideration. Each CCM PDU is approximately 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.

The following section describes some basic hierarchical considerations and the software requirements and configurations that need to be met when considering sub-second enabled MEPs.

- Down MEPs only
- Single peer only
- Any MD Level
 - As long as lower MD level MEPs are not CCM or ETH-APS enabled
G.8031 Ethernet-Tunnels enables OpCode39 Linear APS
G.8032 Ethernet-Rings enables OpCode 40 Ring APS
 - As long as lower MD levels MEPs are not receiving ETH-CCM or ETH-APS PDUs, even if they not locally enabled or configured to do so
The reception of the lower MD level ETH-CCM and ETH-APS PDUs will be processed by the sub second CCM enabled MEP, regardless of MD Level
All other ETH-CFM PDUs will be handled by the MEP at the MD level matching the PDU that has arrived, assuming one has been configured
- Service MEPs (excluding Primary VLAN MEPs)
 - Ethernet SAPs configured on Port with any Ethernet Encapsulation (null, dot1q or QinQ)
- Facility MEPs
 - Ethernet Port Based MEPs
 - Ethernet LAG Based MEPs
 - Ethernet QinQ Tunnel based MEPs (LAG+VLAN, PORT+VLAN)
 - Base Router IP Interfaces
- Service MEPs and Facility MEPs can simultaneously execute sub second CCM enabled MEPs as these are considered different MEP families.

- General processing rules for Service MEPs and Facility MEPs must be met regardless of the CCM interval. These are included here because of the impact misunderstanding could have on the CCM extraction.
 - All the above rules apply
 - MD level hierarchy must be ensured across different families
 - Facility MEPs are the first processing routine for ETH-CFM PDUs
 - VLAN encapsulation uniqueness must exist when processing the ETH-CFM PDU across the two families

Unique Example: An Ethernet Port Based Facility Down MEP configured on port 1/1/1 and Service Down MEP SAP 1/1/1:100 (dot1q encaps) are unique

Conflict Example: An Ethernet Port Based Facility Down MEP configured on port 1/1/1 and Service Down MEP SAP 1/1/1 (null encaps) are in conflict and cannot coexist. All ETH-CFM PDUs will arrive untagged and the Facility MEP takes precedence.
- G.8031 (Ethernet-Tunnels) support both sub second and 1 second CCM intervals and optionally no CCM. When the MEP is created on a G.8031 Ethernet-Tunnel no other MEP that is any way connected to the G.8031 Ethernet-Tunnel can execute sub second CCM intervals.
 - Facility MEPs are not supported in conjunction with G.8031 (Ethernet-Tunnel MEPs)
- G.8032 (Ethernet-Ring) support both sub second and 1 second CCM intervals and optionally no CCM.
 - Facility MEPs are supported in combination with G.8032 MEPs. However, facility MEPs and G.8032 MEPs cannot both execute sub second CCM where the infrastructure is shared. If the operator configures this combination the last updated sub second MEP will overwrite the previous sub second MEP and interrupt the previous configured MEP causing a defRemoteCCM condition.
- The size of the CCM PDU may be increased by configuring the optional Data TLV. This is accomplished by configuring the ccm-padding-size under the specific MEP. The configured value represents the total length of the Data TLV that will be included with the other CCM PDU informational elements. The **no** form of this command removes the optional Data TLV from the CCM PDU. The operator must consider a CCM PDU is 83 byte size in length (75 base elements plus 8 bytes for port status and interface status). If the size of the optional TLV combined with the size of the CCM PDU exceeds 1500 bytes the packet will be dropped if the MTU is 1518/1522. Applies to the 7450 ESS and 7750 SR only.
- CCM will declare a defect when:
 - it stops hearing from one of the remote MEPs for 3.5 times CC interval
 - it hears from a MEP with a LOWER MD level

-
- it hears from a MEP that is not part of the local MEPs MA
 - it hears from a MEP that is in the same MA but not in the configured MEP list
 - it 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.
 - SenderID TLV may optionally be configured to carry the ChassisID. When configured, this information will be included in CCM messages.
 - Only the ChassisID portion of the TLV will be included.
 - The Management Domain and Management Address fields are not supported on transmission.
 - SenderID TLV is not supported with sub second CCM enabled MEPs.
 - Supported for both service (id-permission) and facility MEPs (facility-id-permission).
 - Alarm notification alarm and reset times are configurable under the MEP. By default, the alarm notification times are set to zero, which means the behavior is immediate logging and resetting. When the value is zero and a previous higher level alarm is reset, if a lower level alarm exist, and is above the low-priority defect, that log event will be created. However, when either of the alarm notification timers are non-zero and a lower priority alarm exists, it will not be logged.
 - Alarm (fng-alarm-time) will delay the generation of the log event by the value configured. The alarm must be present for this amount of time before the log event is created. This is for only log event purposes.
 - Reset (fng-reset-time) is the amount of time the alarm must be absent before it is cleared.

You can use the optional **ccm-tlv-ignore** command to ignore the reception of interface-status and port-status TLVs in the ETH-CCM PDU on Facility MEPs (Port, LAG, QinQ Tunnel and Router). No processing is performed on the ignored ETH-CCM TLVs values.

Any TLV that is ignored is reported as *absent* for that remote peer and the values in the TLV do not have an impact on the ETH-CFM state machine. This the same behavior as if the remote MEP never included the ignored TLVs in the ETH-CCM PDU. If the TLV is not properly formed, the CCM PDU will fail the packet parsing process, which will cause it to be discarded and a defect condition will be raised.

NODE1:

```
config>eth-cfm# info
-----
      domain 4 format none level 4
        association 1 format icc-based name "04-0000000102"
          bridge-identifier 100
          exit
          ccm-interval 1
          remote-mepid 102
        exit
      exit
-----
```

NODE2:

```
config>eth-cfm# info
-----
      domain 4 format none level 4
        association 1 format icc-based name "04-0000000102"
          bridge-identifier 100
          exit
          ccm-interval 1
          remote-mepid 101
        exit
      exit
-----
```

Common CCM attributes are defined within the association, including the list of remote peers and interval. Once this is complete, the MEP configured on the SAP within the service must enabled CCM and the priority of the packet can be set.

NODE1:

```
config>service>epipe# info
-----
      sap 1/1/2:100.31 create
        eth-cfm
          mip mac D0:0D:1E:01:01:01
          exit
        exit
      sap 1/1/10:100.31 create
        eth-cfm
          mep 101 domain 4 association 1 direction up
            ccm-enable
            mac-address d0:0d:1e:00:01:01
            no shutdown
          exit
        exit
-----
```

```

        exit
    exit
    no shutdown
-----

NODE2:

config>service>epipe# info
-----
    sap 1/1/2:100.31 create
        eth-cfm
            mip mac D0:0D:1E:01:01:02
        exit
    exit
    sap 1/1/10:100.31 create
        eth-cfm
            mep 102 domain 4 association 1 direction up
            ccm-enable
            mac-address d0:0d:1e:00:01:02
            no shutdown
        exit
    exit
    exit
    no shutdown
-----

```

There are various display commands that are available to show the status of the MEP and the list of remote peers.

3.4.6 CC Remote Peer Auto-Discovery

As specified in the section “Continuity Checking (CC),” all remote MEP-IDs must be configured under the association using the **remote-mepid** command in order to accept them as peers. When a CCM is received from a MEP-ID that has not been configured, the “unexpected MEP” will cause the defErrorCCM condition to be raised. The defErrorCCM will be raised for all invalid CC reception conditions.

The auto-mep-discovery option allows for the automatic adding of remote MEP-IDs contained in the received CCM. Once learned, the automatically discovered MEP behave the same as a manually configured entry. This includes the handling and reporting of defect conditions. For example, if an auto discovered MEP is deleted from its host node, it will experience the standard timeout on the node which auto discovered it.

Obviously, when this function is enabled, the “unexpected MEP” condition no longer exists. That is because all MEPs are accepted as peers and automatically added to the MEP database upon reception. There is an exception to this statement. If the maintenance association has reached its maximum MEP count, and no new MEPs can be added, the “unexpected MEP” condition will raise the defErrorCCM defect condition. This is because the MEP was not added to the association and the remote MEP is still transmitting CCM.

The **clear eth-cfm auto-discovered-meps** [*mep-id*] **domain** *md-index* **association** *ma-index* is available to remove auto discovered MEPs from the association. When the optional *mep-id* is included as part of the clear command, only that specific MEP-ID within the domain and association will be cleared. If the optional *mep-id* is omitted when the clear command is issued, all auto discovered MEPs that match the domain and association will be cleared. The clear command is only applicable to auto discovered MEPs.

If there is a failure to add a MEP to the MEP database and the action was manual addition using the “remote-mepid” configuration statement, the error “MINOR: ETH_CFM #1203 Reached maximum number of local and remote endpoints configured for this association” will be produced. When failure to add a MEP to the database through an auto discovery, no event is created. The CCM Last Failure indicator tracks the last CCM error condition. The decode can be viewed using the “show eth-cfm mep *mep-id* domain *md-index* association *ma-index*” command. An association may include both the manual addition of remote peers using the remote-mepid and the auto-mep-discovery option.

The all-remote-mepid display includes an additional column AD to indicate where a MEP has been auto discovered, using the indicator T.

Auto discovered MEPs will not survive a system reboot. These are not permanent additions to the MEP database and will be not reload after a reboot. The entries will be relearned when the CCM is received. Auto discovered MEPs can be changed to manually created entries simply by adding the appropriate remote-mepid statement to the proper association. At that point, the MEP is no longer considered auto discovered and can no longer be cleared.

If a remote-mepid statement is removed from the association context and auto-mep-discovery is configured and a CC message arrives from that remote MEP, it will be added to the MEP database, this time as an auto discovered MEP.

The individual MEP database for an association must not exceed the maximum number of MEPs allowed. A MEP database consists of all local MEPs plus all configured remote-mepids and all auto discovered MEPs. If the number of MEPs in the association has reached capacity, no new MEPs may be added in any manner. The number of MEPs must be brought below the maximum value before MEPs can be added. Further, the number of MEPs across all MEP databases must not exceed the system maximum. The number of MEPs supported per association and the total number of MEPs across all associations is dependent of the system SF/CPM.

3.4.7 CCM Grace Period

When an ISSU operation or soft reset function is invoked, the ETH-Vendor Specific Message (ETH-VSM) PDU is used to announce a grace period to a remote CCM enabled peer which are administratively enabled. This Multicast Class 1 DA announcement includes the start of a grace period, the new remote timeout value of 90s and the completion of the grace process. Those MEPs configured with unicast destination MAC addresses will still receive the CCM messages as unicast.

At the start of the operation, a burst of three packets will be sent over a three second window in order to reduce the chances that a remote peer may miss the backoff announcement. This grace announcement will include an indication that the local node that is undergoing a maintenance operation that could possibly delay the announcement of CCM messages at the configured interval.

Three evenly spaced ETH-VSM messages will be sent during the interval advertised in the ETH-VSM message. This means that the ETH-VSM message will be sent every 10 seconds to all appropriate remote peers. Reception of this packet refreshes the timeout calculation. The local node undergoing the maintenance operation will also delay the CCM timeout by the announced ETH-VSM interval. This local interval will be reset when any ETH-CC PDU is received on the MEP. An optional TLV is included in AIS packets to extend timeout values for active AIS conditions.

At the end of the maintenance operation there will be a burst of three more messages over a 10 second window that will indicate that the maintenance operation has completed. Once the first of these messages has been received the receiving peer will transition back to the ETH-CCM message and associated interval as the indication for the remote timeout ($3.5 * \text{ccm-interval} + \text{hold if any}$).

CCM message will continue to be sent during this process but loss of the CCM packets during this 10s window will not affect the remote peer timeout. The only change to the CCM processing is which timer to use during the maintenance operation. During the operation, the value used is that announced as part of the ETH-VSM message. Outside a maintenance window the standard $\text{CCM-interval} \times 3.5$ + any configured hold time is used. Since CCM messages are sent during this time other faults and failures can still be conveyed and acted upon. These include AIS, Interface status settings, etc. Only the remote peer timeout (defRemoteCCM) is affected by the ETH-VSM announcement.

The grace announcement using ETH-VSM will continue until the upgrade or reset is completed. During an IOM or line card soft reset ETH-CFM will not determine which peers are affected by a soft reset of a specific IOM or line card. All remote peers will receive the ETH-VSM with the grace period announcement until the soft reset is completed. This means that all remote MEPs, regardless of location on the local node will enter a grace.

Clearing the IOM or line card does not invoke the organizational specific TLV with the grace period announcement.

This is a value added function that is applicable to only nodes that implement support for ALU's approach for announcing grace using ETH-VSM. As specified in the standards, when a node does not support a specific optional function the message will be ignored and the no processing will be performed.

This feature is enabled by default. A system wide command is available to disable this transmission of these grace messages. Entering the no grace-tx-enable in the configuration under the **eth-cfm>system** context will prevent the grace announcements. If this configuration option is change from enable to disable while grace is being announced the three grace stop messages will be transmitted. Changing the state of this configuration option from disable to enable will only affect future ISSU and soft reset functions. It will have no affect on any ISSU or soft reset function that is active at the time this command was enabled.

3.4.8 CCM Hold Timers

In some cases the requirement exists to prevent a MEP from entering the defRemoteCCM defect, remote peer timeout, from more time than the standard 3.5 times the CCM-interval. Both the IEEE 802.1ag standard and ITU-T Y.1731 recommendation provide a non-configurable 3.5 times the CCM interval to determine a peer time out. However, when sub second CCM timers (10ms/100ms) are enabled the carrier may want to provide additional time for different network segments to converge before declaring a peer lost because of a timeout. In order to maintain compliance with the specifications the ccm-hold-timer down <delay-down> option

has been introduced to artificially increase the amount of time it takes for a MEP to enter a failed state should the peer time out. This timer is only additive to CCM timeout conditions. All other CCM defect conditions, like defMACStatus, defXconCCM, and so on, will maintain their existing behavior of transitioning the MEP to a failed state and raising the proper defect condition without delay.

When the **ccm-hold-timer down** *delay-down* option is configured the following calculation is used to determine the remote peer time out (3.5 times the CCM-Interval + ccm-hold-timer delay-down).

This command is configured under the association. Only sub second CCM enabled MEPs support this hold timer. Ethernet-Tunnel Paths use a similar but slightly different approach and will continue to utilize the existing method. Ethernet-tunnels will be blocked from using this new hold timer.

It is possible to change this command on the fly without deleting it first. Simply entering the command with the new values will change to values without having to delete the command prior to the change.

It is possible to change the ccm-interval of a MEP on the fly without first deleting it. This means it is possible to change a sub second CCM enabled MEP to 1 second or above. The operator will be prevented from changing an association from a sub second CCM interval to a non-sub second CCM interval when a ccm-hold-timer is configured in that association. The ccm-hold-timer must be removed using the no option prior to allowing the transition from sub second to non-sub second CCM interval.

3.4.9 Alarm Indication Signal (ETH-AIS Y.1731)

Alarm Indication Signal (AIS) provides a Y.1731-capable MEP with 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 that is 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 as 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.

AIS generation is not subject to the CCM low-priority-defect parameter setting. When enabled, AIS is generated if the MEP enters any defect condition, by default this includes CCM RDI condition.

To prevent the generation of AIS for the CCM RDI condition, the AIS version of the low-priority-defect parameter (under the **ais-enable** command) can be configured to ignore RDI by setting the parameter value to `macRemErrXcon`. The low-priority-defect parameter is specific and influences the protocol under which it is configured. When the low-priority-defect parameter is configured under CCM, it only influences CCM and not AIS. When the low-priority-defect parameter is configured under AIS, it only influences AIS and not CCM. Each protocol can make use of this parameter using different values.

AIS configuration has two components: receive and transmit. AIS reception is enabled when the command **ais-enable** is configured under the MEP. The transmit function is enabled when the **client-meg-level** is configured.

Alarm Indication Signal function is used to suppress alarms at the client (sub) layer following detection of defect conditions at the server (sub) layer. Due to independent restoration capabilities provided within the Spanning Tree Protocol (STP) environments, ETH-AIS is not expected to be applied in the STP environment.

Transmission of frames with ETH-AIS information can be enabled or disabled on a MEP. Frames with ETH-AIS information can be issued at the client MEG Level by a MEP, including a Server MEP, upon detecting the following conditions:

- Signal failure conditions in the case that ETH-CC is enabled.
- AIS condition in the case that ETH-CC is disabled.

For a point-to-point ETH connection at the client (sub) layer, a client layer MEP can determine that the server (sub) layer entity providing connectivity to its peer MEP has encountered defect condition upon receiving a frame with ETH-AIS information. Alarm suppression is straightforward since a MEP is expected to suppress defect conditions associated only with its peer MEP.

For multipoint ETH connectivity at the client (sub) layer, a client (sub) layer MEP cannot determine the specific server (sub) layer entity that has encountered defect conditions upon receiving a frame with ETH-AIS information. More importantly, it cannot determine the associated subset of its peer MEPs for which it should suppress alarms since the received ETH-AIS information does not contain that information. Therefore, upon receiving a frame with ETH-AIS information, the MEP will suppress alarms for all peer MEPs whether or not there is still connectivity.

Only a MEP, including a Server MEP, is configured to issue frames with ETH-AIS information. Upon detecting a defect condition the MEP can immediately start transmitting periodic frames with ETH-AIS information at a configured client MEG Level. A MEP continues to transmit periodic frames with ETH-AIS information until the defect condition is removed. Upon receiving a frame with ETH-AIS information from its server (sub) layer, a client (sub) layer MEP detects AIS condition and suppresses alarms associated with all its peer MEPs. A MEP resumes alarm generation upon detecting defect conditions once AIS condition is cleared.

AIS may also be triggered or cleared based on the state of the entity over which it has been enabled. Including the optional command **interface-support-enable** under the **ais-enable** command will track the state of the entity and invoke the appropriate AIS action. This means that operators are not required to enable CCM on a MEP in order to generate AIS if the only requirement is to track the local entity. If a CCM enabled MEP is enabled in addition to this function then both will be used to act upon the AIS function. When both CCM and interface support are enabled, a fault in either will trigger AIS. In order to clear the AIS state, the entity must be in an UP operational state and there must be no defects associated with the MEP. The interface support function is available on both service MEPs and facility MEPs both in the Down direction only, with the following exception. An Ethernet QinQ Tunnel Facility MEP does not support interface-support-enable. Many operational models for Ethernet QinQ Tunnel Facility MEPs are deployed with the SAP in the shutdown state.

The following specific configuration information is used by a MEP to support ETH-AIS:

- Client MEG Level — MEG level at which the most immediate client layer MIPs and MEPs exist.
- ETH-AIS transmission period — Determines the transmission period of frames with ETH-AIS information.
- Priority — Identifies the priority of frames with ETH-AIS information.
- Drop Eligibility — Frames with ETH-AIS information are always marked as drop ineligible.
- Interface-support-enable — Optional configuration to track the state of the entity over which the MEP is configured.
- Low-priority-defect — Optional configuration to exclude the CCM RDI condition from triggering the generation of AIS.

A MIP is transparent to frames with ETH-AIS information and therefore does not require any information to support ETH-AIS functionality.

It is important to note that Facility MEPs do not support the generation of AIS to an explicitly configured endpoint. An explicitly configured endpoint is an object that contains multiple individual endpoints, as in pseudowire redundancy.

AIS is enabled under the service and has two parts, receive and transmit. Both components have their own configuration option. The **ais-enable** command under the SAP allows for the processing of received AIS packets at the MEP level. The **client-meg-level** command is the transmit portion that generates AIS if the MEP enter a fault state.

```
config>service>epipe# info
-----
      sap 1/1/2:100.31 create
      eth-cfm
      mip mac D0:0D:1E:01:01:01
      exit
    exit
  sap 1/1/10:100.31 create
  eth-cfm
    mep 101 domain 4 association 1 direction up
    ais-enable
    client-meg-level 5
    exit
    ccm-enable
    mac-address d0:0d:1e:00:01:01
    no shutdown
  exit
exit
exit
no shutdown
-----
```

When MEP 101 enters a defect state, it starts to generate AIS out the passive side of the MEP, away from the fault. In this case, the AIS generates out sap 1/1/10:100.31 since MEP 101 is an up MEP on that SAP. The **Defect Flag** indicates that an RDI error state has been encountered. The **Eth-Ais Tx Counted** value is increasing, indicating that AIS is actively being sent.

A single network event may, in turn, cause the number of AIS transmissions to exceed the AIS transmit rate of the network element. A pacing mechanism is in place to assist the network element to gracefully handle this overload condition. Should an event occur that causes the AIS transmit requirements to exceed the AIS transmit resources, a credit system is used to grant access to the resources. Once all the credits have been used, any remaining MEPs attempting to allocate a transmit resource will be placed on a wait list, unable to transmit AIS. Should a credit be released, when the condition that caused the MEP to transmit AIS is cleared, a MEP on the wait list will consume the newly available credit. If it is critical that AIS transmit resources be available for every potential event, consideration must be given to the worst case scenario and the configuration should never exceed the potential. Access to the resources and the wait list are ordered and maintained in first come first serve basis.

A MEP that is on the wait list will only increment the “Eth-Ais Tx Fail” counter and not the “Eth-Ais TxCount” for every failed attempt while the MEP is on the wait list.

There is no synchronization of AIS transmission state between peer nodes. This is particularly important when AIS is used to propagate fault in ETH-CFM MC-LAG linked designs.

3.4.10 Client Signal Fail (ETH-CSF Y.1731)

Client signal fail (CSF) is a method that allows for the propagation of a fault condition to a MEP peer, without requiring ETH-CC or ETH-AIS. The message is sent when a MEP detects an issue with the entity in the direction the MEP to its peer MEP. A typical deployment model is an UP MEP configured on the entity that is not executing ETH-CC with its peer. When the entity over which the MEP is configured fails, the MEP can send the ETH-CSF fault message.

In order to process the reception of the ETH-CSF message, the **csf-enable** function must be enabled under the MEP. When processing of the received CSF message is enabled, the CSF is used as another method to trigger fault propagation, assuming fault propagation is enabled. If CSF is enabled but fault propagation is not enabled, the MEP will show state of CSF being received from the peer. And lastly, when there is no fault condition, the CSF Rx State will display DCI (Client defect clear) indicating there are no existing failures, even if no CSF has been received. The CSF Rx State will indicate the various fault and clear conditions received from the peer during the event.

CSF carries the type of defect that has been detected by the local MEP generating the CSF message.

- 000 – LOS – Client Loss of Signal
- 001 – FDI/AIS – Client forward defect indication
- 010 – RDI – Client reverse defect indication

Clearing the CSF state can be either implicit, time out, or explicit, requiring the client to send the PDU with the clear indicator (011 – DCI – Client defect clear indication). The receiving node uses the multiplier option to determine how to clear the CSF condition. When the multiplier is configured as non zero (in increments of half seconds between 2 and 30) the CSF will be cleared when CSF PDUs have not been received for that duration. A multiplier value of 0 means that the peer that has generated the CSF must send the 011 – DCI flags. There is no timeout condition.

Service-based MEP supports the reception of the ETH-CSF as an additional trigger for the fault propagation process. Primary VLAN and Virtual MEPs do not support the processing of the CSF PDU. CSF is transparent to MIPs. There is no support for the transmission of ETH-CSF packets on any MEP.

3.4.11 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.

Specific configuration information required by a MEP to support ETH-test is the following:

- MEG level — MEG level at which the MEP exists
- Unicast MAC address of the peer MEP for which ETH-test is intended.
- Data - Optional element whose length and contents are configurable at the MEP.
- Priority — Identifies the priority of frames with ETH-Test information.
- Drop Eligibility — Identifies the eligibility of frames with ETHTest information to be dropped when congestion conditions are encountered.

A MIP is transparent to the frames with ETH-Test information and does not require any configuration information to support ETH-Test functionality.

Both nodes require the eth-test function to be enabled in order to successfully execute the test. Since this is a dual-ended test, initiate on sender with results calculated on the receiver, both nodes need to be check to see the results.

```
NODE1
config>service>epipe# info
-----
      sap 1/1/2:100.31 create
      eth-cfm
      mip mac D0:0D:1E:01:01:01
      exit
    exit
  sap 1/1/10:100.31 create
  eth-cfm
    mep 101 domain 4 association 1 direction up
    eth-test-enable
    exit
    mac-address d0:0d:1e:00:01:01
    no shutdown
    exit
  exit
exit
no shutdown
-----
# oam eth-cfm eth-test d0:0d:1e:00:01:02 mep 101 domain 4 association 1 data-
length 1000
# oam eth-cfm eth-test d0:0d:1e:00:01:02 mep 101 domain 4 association 1 data-
length 1000
# oam eth-cfm eth-test d0:0d:1e:00:01:02 mep 101 domain 4 association 1 data-
```

```

length 1000

NODE2
config>service>epipe# info
-----
      sap 1/1/2:100.31 create
      eth-cfm
      mip mac D0:0D:1E:01:01:02
      exit
    exit
  sap 1/1/10:100.31 create
  eth-cfm
  mep 102 domain 4 association 1 direction up
  eth-test-enable
  exit
  mac-address d0:0d:1e:00:01:02
  no shutdown
  exit
    exit
  exit
  no shutdown
-----

# show eth-cfm mep 102 domain 4 association 1 eth-test-results
=====
Eth CFM ETH-Test Result Table
=====

```

Peer Mac Addr	FrameCount ByteCount	Current ErrBits CrcErrs	Accumulate ErrBits CrcErrs
d0:0d:1e:00:01:01	3	0	0
	3000	0	0

```

=====

```

3.4.12 One-Way Delay Measurement (ETH-1DM Y.1731)

One-way delay measurement allows the operator the ability to check unidirectional delay between MEPs and applies to the 7450 ESS and 7750 SR only. An ETH-1DM packet is time stamped by the generating MEP and sent to the remote node. The remote node time stamps the packet on receipt and generates the results. The results, available from the receiving MEP, will indicate the delay and jitter. Jitter, or delay variation, is the difference in delay between tests. This means the delay variation on the first test will not be valid. It is important to ensure that the clocks are synchronized on both nodes to ensure the results are accurate. NTP can be used to achieve a level of wall clock synchronization between the nodes.



Note: Accuracy relies on the nodes ability to timestamp the packet in hardware. Network elements that do not support this hardware time stamping, like the ESS-1, will display different results than hardware timestamp capable devices, like the SR-7/SR-12 and ESS-7/ESS-12.

3.4.13 Two-Way Delay Measurement (ETH-DMM Y.1731)

Two-way delay measurement is similar to one way delay measurement except it measures the round trip delay from the generating MEP. In this case wall clock synchronization issues will not influence the test results because four timestamps are used. This allows the remote nodes time to be removed from the calculation and as a result clock variances are not included in the results. The same consideration for first test and hardware based time stamping stated for one way delay measurement are applicable to two-way delay measurement.

Delay can be measured using one-way and two-way on demand functions. The two-way test results are available single-ended, test initiated, calculation and results viewed on the same node. There is no specific configuration under the MEP on the SAP in order to enabled this function. An example of an on demand test and results are below. The latest test result is stored for viewing. Further tests will overwrite the previous results. Delay Variation is only valid if more than one test has been executed.

```
oam eth-cfm two-way-delay-test d0:0d:1e:00:01:02 mep 101 domain 4 association 1
```

Two-Way-Delay-Test Response:

Delay 2955 microseconds Variation 111 microseconds

```
# show eth-cfm mep 101 domain 4 association 1 two-way-delay-test
```

```
=====
Eth CFM Two-way Delay Test Result Table
=====
Peer Mac Addr          Delay (us)          Delay Variation (us)
-----
d0:0d:1e:00:01:02      2955                111
=====
```

3.4.14 Synthetic Loss Measurement (ETH-SLM Y.1731)



Note: Release 9.0R1 uses pre-standard OpCodes and will not interoperate with any other release or future release.

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. ALU 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/are lost, the count + unacknowledged equals the number of probes sent.
- Out-Loss (Far-end) — Packets lost on the way to the remote node, from test initiator to 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 means 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 up to 1000 seconds in length. Although possible, it is more likely the majority of on demand case will be run up 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 will 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/jitter with a single test. The implementation does not cause any interoperability because the optional TVL will be ignored by equipment that does not support this. In mixed vendor environments loss measurement will continue to be tracked but delay and jitter will only report round trip times. It is important to point out that the round trip times in this mixed vendor environments will 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 will be reported. Since all four time stamps are included in the packet the round trip time in this case will 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. In order to completely understand how SAA functions please refer to the appropriate section of the user guide.

The ETH-SL packet format contains a test-id that will be internally generated and not configurable. The test-id will be visible for the on demand test in the display summary. It is possible a remote node processing the SLM frames will receive 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 will cause 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 somewhat likely. In this release, only the first responder will be used to measure packet loss. The second responder will be 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 value, the responding node releases the index in the table and responds with a sequence number of 1, regardless of the sequence number sent by the instantiating node. Expiration of this timer causes the reflecting peer to expire the previous test. Any packets that follow the expiration of a test will be viewed as a new test. The default for the inactivity-timer is 100 second and has a range of ten to 100 seconds.


Only the configuration is supported by HA. There will be no synchronization of data between active and standby. Any unwritten, or active tests will be lost during a switchover and the data will not be recoverable.

ETH-SL provides a mechanism for operators to pro-actively trend packet loss.

3.4.14.1 Configuration Example

Figure 37 shows the configuration required for proactive SLM test using SAA.

Figure 37 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
-----

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
-----

config>saa# info
-----
      test "slm1"
        type
          eth-cfm-two-way-slm d0:0d:1e:00:01:01 mep 100 domain 3
        association 1 count 100 timeout 1 interval 1
        exit
        continuous
        no shutdown
      exit
-----
```

The following sample output is meant to demonstrate the different loss conditions that an operator may see. The total number of attempts is “99” is because the final probe in the test was not acknowledged.

```
# show saa slm1
Test Run: 183
Total number of attempts: 99
```

```

Number of requests that failed to be sent out: 0
Number of responses that were received: 48
Number of requests that did not receive any response: 50
Total number of failures: 50, Percentage: 50
(in ms)           Min           Max           Average           Jitter
Outbound  :       -370          -362           -366            0.432
Inbound   :         363           371            367            0.308
Roundtrip :         0.000         5.93            1.38            0.496
Per test packet:
Sequence Outbound Inbound RoundTrip Result
    1      0.000    0.000    0.000 Out Loss
    2      0.000    0.000    0.000 Out Loss
    3      0.000    0.000    0.000 Out Loss
    4      0.000    0.000    0.000 Out Loss
...snip...
   46     -369      370      1.28 Response Received
   47     -362      363      1.42 Response Received
   48      0.000    0.000    0.000 In Loss
   49      0.000    0.000    0.000 In Loss
   50     -362      363      1.42 Response Received
   51     -362      363      1.16 Response Received
   52     -362      364      1.20 Response Received
   53     -362      364      1.18 Response Received
   54     -363      364      1.20 Response Received
...snip...
   96     -369      370      1.29 Response Received
   97     -369      370      1.30 Response Received
   98      0.000    0.000    0.000 Unacknowledged
   99      0.000    0.000    0.000 Unacknowledged
  100      0.000    0.000    0.000 Unacknowledged

```

```
=====
```

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 d0:0d:1e:00:01:01 mep 100 domain 3 association 1 send-
count
20 interval 1 timeout 1

```

```

Sending 20 packets to d0:0d:1e:00:01:01 from MEP 100/3/1 (Test-id: 588)

```

```

Sent 20 packets, 20 packets received from MEP ID 101, (Test-id: 588)
(0 out-loss, 0 in-loss, 0 unacknowledged)

```

```

# show eth-cfm mep 100 domain 3 association 1 two-way-slm-test

```

```

=====
Eth CFM Two-way SLM Test Result Table (Test-id: 588)
=====

```

Peer Mac Addr	Remote MEP	Count	In Loss	Out Loss	Unack
d0:0d:1e:00:01:01	101	20	0	0	0

```

=====

```

3.4.15 Frame Loss Measurement (ETH-LMM Y.1731)

The ETH-LMM Y.1731 approach to Ethernet loss measurement allows for the collection of frame counters in order to determine the unidirectional frame loss between point-to-point ETH-CFM MEP peers. This loss measurement does not rely on the counting of its own PDU in order to determine unidirectional loss. The protocol PDU includes four counters which represent the data sent and received in each direction: Transmit Forward (TxFCf), Receive Forward (RxFCf), Transmit Backward (TxFCb) and the Receive Backward (RxFC1).

The protocol is designed only for point-to-point connections. It is impossible for the protocol to properly report loss if the point-to-point relationship is broken; for example, if a SAP or MPLS binding is receiving data from multiple peers, as could be the case in VPLS deployments, this protocol cannot be used in any reliable fashion.

The loss differential between transmit and receive is determined the first time an LMM PDU is sent. Each subsequent PDU for a specific test will perform a computation of differential loss from that epoch. Each processing cycle for an LMR PDU will determine if there is a new maximum of minimum loss window, add any new loss to the frame loss ratio computation, and update the four raw transmit and receive counters. The individual probe results are not maintained; these results are only used to determine a new minimum of maximum. A running total of all Tx and Rx values is used to determine the average Frame Loss Ratio (FLR) at the completion of the measurement interval. The data set includes the protocol information in the opening header, followed by the frame counts in each direction, and finally the FLR percentages.


Service frame recording does have some caveats that need to be understood before selecting this method of loss measurement. Statistics are maintained per forwarding complex. Multiple path environments may spread frames between the same two peers across different forwarding complexes (for example, link aggregation groups). The ETH-LMM Y.1731 protocol has no means of rationalizing different transmit and receive statistics when there are complex changes or when any statistics have been cleared on either of the peer entities. The protocol will resynchronize but the data collected for that measurement interval will be invalid. The protocol has no method to determine if the loss is true loss or whether some type of complex switch has occurred or statistics were cleared and as such cannot use any suspect flag to mark the data as invalid. Higher level systems must coordinate network events and administrative actions that can cause the counters to become non-representative of the service data loss.

Packet reordering also affect frame loss and gain reporting. If there is queuing contention on the local node, or if there are path differences in the network that cause frames to be interleaved or delayed, the counter stamped into the LMM PDU could introduce frame gain or loss in either direction. For example, if the LMM PDU is stamped with the TxFCf counter and the LMM PDU traffic is interleaved but the interleaving can not be accounted for in the counter, then a potential gain would be realized in the forward direction. This is because the original counter included as the TxFCf value would not have included those interleaved packets, but the RxFCf counter on the remote peer would include those packets. Gains and losses will even out over the life of the measurement interval. Absolute values will be used for any negative values, per interval or at the end of the measurement interval.

With ETH-LMM Y.1731, a single per SAP or per MPLS SDP binding or per facility counter is maintained. This single counter model applies to any conflicting entity that attempts to collect per entity statistics that may cover the same resource. This means that per service and facility MEP LMM counting is not supported. The operator must choose one type of facility MEP or the service level MEP. If a facility MEP is chosen (Port, LAG, QinQ Tunnel or Base Router Interface) care must be taken to ensure the highest configured MEP performs the loss collection routine. Configuring loss collection on a lower level MEP will lead to additive gain introduced in both directions. Although the collection statement is not blocked by CLI or SNMP when there are potential conflicts only one will be accurate. The operator must be aware of lower level resource conflicts. For example, a null based service SAP, any default SAP context or SAP that covers the entire port or facility resource, such as sap 1/1/1, will always count the frame base loss counter against the SAP and never the port, regardless of the presences of a MEP or the **collect-lmm-stats** configuration on the SAP. Resource contention extends beyond the sharing of common resources used for packet counting and extraction. There is also protocol level contention. For example, the Cflowd cannot be counted or sampled on an entity that is collecting LMM stats. Collection of per Ethernet SAP or per MPLS SDP binding or per facility is not enabled by default. In order for this feature to function with accurate measurements, the **collect-lmm-stats** is required under the ETH-CFM context for the Ethernet SAP or MPLS SDP binding or under the MEP in the case of the facility MEP. If this command is not enabled on the launch or reflector, the data in the LMM and LMR PDU will not be representative and the data captured will be invalid. The **show>service>sdp-using eth-cfm** and **show>service>sap-using eth-cfm** commands have been expanded to include the **collect-lmm-stats** option for service based MEPs. The **show>eth-cfm>cfm-stack-table facility** command has been expanded to include **collect-lmm-stats** to view all facility MEPs. Using these commands with this new option will display which entities are currently collecting LMM counter.

The counter will include all frames that are transmitted or received regardless of class of service or discard eligibility markings. Locally transmitted and locally terminated ETH-CFM frames on the peer collecting the statistics will not be included in the counter. However, there are deployment models that will introduce artificial frame loss or gain when the ETH-CFM launch node and the terminating node for some ETH-CFM packets are not the same peers. [Figure 38](#) demonstrates this issue.

Figure 38 Mismatched LMM Statistical Counters



Launching a single-ended test is under the control of the OAM Performance Monitoring (OAM-PM) architecture and as such adheres to those rules. The ETH-LMM Y.1731 functionality is not available through interactive CLI or interactive SNMP, it is only available under the OAM-PM configuration. This includes the assignment of a Test-Id. This protocol does not carry the 4-byte test id in the packet. This is for local significance and uniformity with other protocols under the control of the OAM-PM architecture. A single LMM test can be configured and executed between the same two ETH-CFM MEP peers. Support is included for point-to-point up and Up and Down Service MEPs and Down Facility MEPs (Port, LAG and Base Router Interfaces). Base router interface accuracy may be affected by the layer two layer three inter-working functions, routing protocol, ACLs, policies, and other layer three functions that were never meant to be accounted for by an Ethernet frame loss measurement tool. Launch functions require IOM3/IMM and beyond as well as a minimum of SF/CPM3. A node reflecting the ETH-LMM PDU requires IOM3/IMM but does not require a SF/CPM3.

There is no support for ETH-LMM UP MEPs in an I-VPLS or PBB ePipe. Configuring this will result in LMM PDU being discarded on the remote BVPLS peer. There is no support for ETH-LMM when Primary VLANs are configured against the MEP. If the SAP over which the UP MEP is configured is not operational the LMM or LMR transmissions will fail. This is because the SAP which stores the counters is unavailable to the LMM PDU. QinQ Tunnel collection will be the aggregate of all

outer VLANs that share the VLAN with the tunnel. If the QinQ is configured to collect LMM statistics then any Service MEP that shares the same VLAN as the QinQ tunnel will be blocked from configuring the **collect-imm-stats** command. The reverse is also true. If a fully qualified SAP is configured to collect LMM statistics the QinQ tunnel that shares that outer VLAN will be block from configuring **collect-imm-stats**.

3.5 ETH-CFM Statistics

A number of statistics are available to view the current overall processing requirements for CFM. Any packet that is counted against the CFM resource will be included in the statistics counters. These counters do not include the counting of sub-second CCM, ETH-CFM PDUs that are generated by non-ETH-CFM functions (which includes OAM-PM & SAA) or are filtered by an applicable security configuration.

SAA and OAM-PM use standard CFM PDUs. The reception of these packets are included in the receive statistics. However, these two functions are responsible for launching their own test packets and do not consume ETH-CFM transmission resources.

Per system and per MEP statistics are available with a per OpCode breakdown. Use the **show eth-cfm statistics** command to view the statistics at the system level. Use the **show eth-cfm mep mep-id domain md-index association ma-index statistics** command to view the per MEP statistics. These statistics may be cleared by substituting the **clear** command for the **show** command. The clear function will only clear the statistics for that function. For example, clear the system statistics does not clear the individual MEP statistics, each maintain their own unique counters.

All known OpCodes are listed in transmit and receive columns. Different versions for the same OpCode are not distinguished for this display. This does not imply the network element supports all listed functions in the table. Unknown OpCodes will be dropped.

It is also possible to view the top ten active MEPs on the system. The term active can be defined as any MEP that is in a “no shutdown” state. The **tools dump eth-cfm top-active-meps** can be used to see the top ten active MEPs on the system. The counts will be based from the last time to command was issued with the **clear** option. MEPs that are in a shutdown state are still terminating packets, but these will not show up on the active list.

```
tools dump eth-cfm top-active-meps
Calculating most active MEPs in both direction without clear ...
```

MEP	Rx Stats	Tx Stats	Total Stats
12/4/28	3504497	296649	3801146
14/1/28	171544	85775	257319
14/2/28	150942	79990	230932

```
tools dump eth-cfm top-active-meps clear
Calculating most active MEPs in both direction with clear ...
```

MEP	Rx Stats	Tx Stats	Total Stats
12/4/28	3504582	296656	3801238

14/1/28	171558	85782	257340
14/2/28	150949	79997	230946

tools dump eth-cfm top-active-meps clear
Calculating most active MEPs in both direction with clear ...

MEP	Rx Stats	Tx Stats	Total Stats
-----	-----	-----	-----
12/4/28	28	2	30
14/1/28	5	2	7
14/2/28	3	2	5

These statistics help operators to determine the busiest active MEPs on the system as well a breakdown of per OpCode processing at the system and MEP level.

3.6 ETH-CFM Packet Debug

The debug infrastructure supports the decoding of both received and transmitted valid ETH-CFM packets for MEPs and MIPs that have been tagged for decoding. The **eth-cfm** hierarchy has been added to the existing **debug** CLI command tree. When a MEP or MIP is tagged by the debug process, valid ETH-CFM PDUs that allocate ETH-CFM transmit and receive resources will be decoded and presented to the logging infrastructure for operator analysis. Fixed queue limits restrict the overall packet rate for decoding. The receive and transmit ETH-CFM debug queues are serviced independently. Receive and transmit correlation is not maintained across the receive and transmit debug queues. The **tools dump eth-cfm debug-packet** command will display message queue exceptions.

Valid ETH-CFM packets must pass a multiple-phase validity check before being passed to the debug parsing function. MAC addresses must be non-zero. If the destination MAC address is multicast, the last nibble of the multicast address must match the expected level of a MEP or MIP tagged for decoding. Packet length and TLV formation, usage, and, where applicable, field validation is performed. Finally, the OpCode-specific TLV structural checks are performed against the remainder of the PDU.

An ETH-CFM packet that passes the validation process is passed to the debug decoding process for tagged MEPs or MIPs. The decoding process parses the PDU for analysis. Truncation of individual TLVs will occur when:

- TLV processing requires multiple functions; this occurs with TLVs that include sub-fields
- an Organizational Specific TLV exists
- padding has been added, as in the case of the optional Data or Test TLVs
- an unknown OpCode is detected; the decode process will process the generic ETH-CFM header with a hex dump for unknown fields and TLVs

The number of printable bytes is dependent on the reason for truncation.

Any standard fields in the PDU that are defined for a certain length with a Must Be Zero (MBZ) attribute in the specification will be decoded based on the specification field length. There is no assumption that packets adhere to the MBZ requirement in the byte field; for example, the MEP-ID is a 2-byte field with three reserved MBZ bits, which translates into a standard MEP-ID range of 0 to 8191. If the MBZ bits are violated, then the 2-byte field will be decoded using all non-zero bits in the 2-byte field.

The decoding function is logically positioned between ETH-CFM and the forwarding plane. Any ETH-CFM PDU discarded by an applicable security configuration will not be passed to the debug function. Care must be taken when interpreting specific ETH-CFM PDU decodes. Those PDUs that have additional, subsequent, or augmented information applied by the forwarding mechanisms may not be part of the decoded packet. Examples of exclusions include sub-second CCM and ETH-CFM PDUs that are generated by non-ETH-CFM functions (which includes OAM-PM and SAA). Augmentation includes the timestamp applied to ETH-CFM PDUs by the forwarding plane. Any packet that is discarded by a security mechanism, including squelching (**config>service>sap>eth-cfm>squelch-ingress-levels**) or CPU protection (**config>service>sap>eth-cfm>cpu-protection ... eth-cfm-monitoring**), will bypass the decoding function.

This function allows for enhanced troubleshooting for ETH-CFM PDUs to and from tagged MEPs and MIPs. Only defined and node-supported functionality will be decoded, possibly with truncation. Unsupported or unknown functionality on the node is treated on a best-effort basis, typically handled with a decode producing a truncated number of hex bytes.

This functionality does not support decoding of sub-second CCM, any ETH-CFM PDUs that are processed by non-ETH-CFM entities (which includes OAM-PM and SAA CFM transmit functions), or MIPs created using the **default-domain** table.

3.7 ETH-CFM CoS Considerations

UP MEPs and Down MEPs have been aligned 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. The numerical ETH-CFM **priority** value resolves FCs using the following mapping:

- 0 — be
- 1 — l2
- 2 — af
- 3 — l1
- 4 — h2
- 5 — ef
- 6 — h1
- 7 — nc

If there is no egress QoS policy, the priority value will be mapped to the CoS values in the frame. An ETH-CFM frame utilizing VLAN tags will have the DEI bit mark the frame as “discard ineligible”. However, egress QoS Policy may overwrite this original value. The Service Assurance Agent (SAA) uses [fc {fc-name} [profile {in|out}]] 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.

This does not include Ethernet Linktrace Response (ETH-LTR). The specification requires the highest priority on the bridge port should be used in response to an Ethernet Linktrace Message (ETH-LTM). This provides the highest possible chance of the response returning to the source. Operators may configure the linktrace response priority of the MEP using the **ccm-ltm-priority**. MIPs inherit the MEPs priority unless the **mip-ltr-priority** is configured under the bridging instance for the association (**config>eth-cfm>domain>assoc>bridge**).

3.8 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 (E-LMI used for OAM). Another example allows an Ipipe service, where one end is Ethernet and the other end is Frame Relay, ATM, PPP, MLPPP, or HDLC.

In the SR OS 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).

Fault propagation cannot be enabled for eth-tun control MEPs (MEPs configured under the eth-tun primary and protection paths). However, failure of the eth-tun (meaning both paths fail) will be propagated by SMGR because all the SAPs on the eth-tun will go down.

This section describes the following topics:

- [CFM Connectivity Fault Conditions](#)
- [CFM Fault Propagation Methods](#)
- [Epipe Services](#)
- [CFM Detected Fault](#)
- [Ipipe Services](#)
- [VPLS Service](#)
- [IES and VPRN Services](#)
- [Pseudowire Switching](#)
- [LLF and CFM Fault Propagation](#)
- [802.3ah EFM OAM Mapping and Interaction with Service Manager](#)

3.8.1 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: Remote Defect Indication. Remote MEP is declaring a fault by setting the RDI bit in the CCM PDU. Typically a result of raising a local defect based on of the CCM or lack of CCM from an expected or unexpected peer. A feedback loop into the association as a notification since CCM is multicast message with no response.
- DefMACstatus: MAC layer issue. Remote MEP is indicating remote port or interface status not operational.
- DefRemoteCCM: No communication from remote peer. MEP not receiving CCM from an expected remote peer. Timeout of CCM occurs in 3.5 x CC interval.
- DefErrorCCM: Remote configuration does not match local expectations. Receiving CC from remote MEP with inconsistent timers, lower MD/MEG level within same MA/MEG, MEP receiving CCM with its own MEP ID within same MA/MEG.
- DefXconCCM: Cross-connected services. MEP receiving CCM from different MA/MEG.

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. In general terms, when a MEP propagates fault to a peer the peer receiving the fault must not reciprocate with a fault back to the originating MEP with a fault condition equal to or higher than the originating MEP low-priority-defect setting. 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.

3.8.2 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

For using AIS for fault propagation, AIS must be enabled for the MEP. The AIS configuration needs to be updated to support the MD level of the MEP (currently it only supports the levels above the local MD level).

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 is triggered and does not wait for the expiration of the remaining CCM interval transmission. This rule applies to CFM fault notification for all services.

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.

3.8.3 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.

3.8.4 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 operup. 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/SDPbinding; 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.

3.8.4.1 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 and E-LMI) on the mate SAP/SDP-binding.

3.8.4.2 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.

3.8.4.3 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.


3.8.5 Ipipe Services

3.8.5.1 CFM Detected Fault

Deployment of solutions that include legacy to Ethernet aggregation should involve fault inter-working consideration. Protocols like Frame Relay propagate fault using the Local Management Interface (LMI). However, other protocols do not include a dedicated management interface over which to indicate fault. PPP, MLPPP and Cisco HDLC must use a different mechanism to communicate fault between the two different connection types.

The **eth-legacy-fault-notification** option and the associated parameters along with Ethernet CFM fault propagation on the Ethernet SAP MEP must be enabled in order to properly inter-work the Ethernet and PPP, MLPPP or Cisco HDLC connections. [Figure 39](#) shows the various high level functions that inter-work Ethernet aggregation and legacy interfaces using point to point Ipipe services.

Figure 39 Fault Propagation Model



In general the lpipe service requires the ce-address information to be learned or manually configured as part of the Ethernet SAP object before the legacy interface connection can be established. IPv6 includes an optimization that uses the Link Local IPv6 address to start the legacy negotiation process and does not require the ce-addressing described previously. This IPv6 optimization does not align well with fault inter-working functions and is disabled when the **eth-legacy-fault-notification** function is enabled.

Fault propagation is not active from the Ethernet SAP to the legacy connection if the ce-address information for the Ethernet SAP has not been learned or configured. If both IPv4 and IPv6 are configured, each protocol will require ce-addressing to be learned or configured enabling fault inter-working for that protocol. Once the ce-address has been learned or configured for that protocol, fault inter-working will be active for that protocol. If either IPv4 or IPv6 ce-addressing from the Ethernet SAP is resident, the access legacy SAP will be operational. The NCP layer will indicate which unique protocol is operational. Fault propagation toward the Ethernet SAP from the legacy connection will still be propagated even if the ce-address is not resident within the lpipe under the following conditions; if any SAP or the Service is shutdown, or the legacy SAP is not configured.

The learned Ethernet ce-address is a critical component in lpipe service operation and fault propagation. In order to maintain the address information the **keep** option must be configured as part of the **ce-address-discovery** command. If the **keep** command is not configured, the address information is lost when the Ethernet SAP transitions to a non-operational state. When the address information is flushed, the lpipe service will propagate the fault to the legacy PPP, MLPPP and Cisco HDLC connections. The lack of the ce-addressing on the Ethernet SAPs may cause a deadlock condition that requires operator intervention to resolve the issue. The **keep** command must be configured when the **eth-legacy-fault-notification** functionality is enabled with PPP, MLPPP and Cisco HDLC legacy interfaces, and fault

propagation is required using this type of aggregation deployment. The **keep** option is specific to and only supported when **eth-legacy-fault-notification** is configured. If the **keep** option is configured as part of the ce-address-discovery command, the eth-legacy-fault-propagation cannot be removed. Configuration changes to the **ce-address-discovery** command may affect the stored ce-address information. For example, if the eth-legacy-fault-notification **ipv6 keep** is changed to **ce-address-discovery keep**, the stored IPv6 ce-address information is flushed. If the **keep** option is removed, all discovered ce-address information is flushed if the SAP is operationally down.

The ce-address stored in the lpipe service as part of the discovery process will be updated if a new ARP arrives from the layer three device connected to the Ethernet SAP. If the layer three device connected to the Ethernet SAP does not send an ARP to indicate the addressing information has been changed, the ce-address stored locally as part of the previous discovery function will be maintained. If changes are made to the layer three device connected to the Ethernet SAP that would alter the ARP information and that device does not generate an ARP packet, or the lpipe inter-working device does not receive the ARP packet, for example, the Ethernet SAP is admin down for IPv4, or the service is operationally down for IPv6, the stored ce-address retained by the lpipe as a result of the keep operation will be stale. This stale information will result in a black hole for service traffic. The **clear service id service-id arp** can be used to flush stale ARP information. This will not solicit a arp from a peer.

The **keep** option will not maintain the ce-address information when the Ethernet SAP is administratively shutdown or when the node reboots.

Once all the ce-addressing has been populated in the lpipe the legacy interfaces establishment will commence. The successful establishment of these connections will render the lpipe service functional. Legacy connection faults and Ethernet SAP faults may now be propagated.


Should the Ethernet SAP enter a non-operational state as a result of a cable or validation protocol (ETH-CCM), the fault will be inter-worked with the specific legacy protocol. Ethernet faults will inter-work with the legacy interfaces in the following manner:

- PPP: LCP and all NCPs will be shutdown and a terminate-request sent to the far-end.
- MLPPP: LCP will remain operational but the NCP will be shutdown
- HDLC: Suspension of the keepalive messages. The keepalive interval will influence the recovery time. If the recovery timer (discussed later) is equal to the keepalive interval, recovery of the legacy interface recovery may occur after a fault is propagated toward the Ethernet network.

- Frame Relay (does not include support for the **eth-legacy-fault-propagation**):
Signal using LMI messaging

As previously stated, inter-working faults on the legacy connection with the Ethernet infrastructure requires a Down MEP with CCM-enabled configured on the Ethernet SAP with fault-propagation enabled. There are two different methods to propagate fault from a CCM-enabled MEP; **use-int-tlv** or **suspend-ccm**. The **use-int-tlv** approach will cause the CCM message to include the Interface Status TLV with a value of is Down. This will raise a defMACStatus priority error on the peer MEP. The **suspend-ccm** approach will cause the local MEP to suspend transmissions of the CCM messages to the peer MEP. This will raise a defRemoteCCM timeout condition on the peer. The peer must accept these notifications and processes these fault conditions on the local MEP. When the MEP receives these errors, it must not include a defect condition in the CCM messages it generates that is above the peers **low-priority-defect** setting. In standard operation, the MEP receiving the error should only set the RDI bit in the CCM header. If the MEP improperly responds with a defect condition that is higher than the low-priority-defect of the MEP that had generated the initial fault then a deadlock condition will occur and operator intervention will be required. The two CFM propagation methods and the proper responses are shown in [Figure 40](#).

Figure 40 Fault Propagation from Legacy to Ethernet



From a protocol (NCP) perspective, PPP and MLPPP connections have a micro view. Those connections understand the different protocols carried over the PPP and MLPPP connections, and individual protocol errors that can occur. The Ethernet SAP has a macro view without this layer three understanding. When the dual stack IPv4 and IPv6 is deployed, fault can only be propagated from the legacy connection toward the associated Ethernet SAP if both protocols fail on the PPP or MLPPP. If either of the protocols are operational then PPP or MLPPP will not propagate fault in the direction of the Ethernet connection.

Ethernet connection faults are prioritized over legacy faults. When an Ethernet fault is detected, any fault previously propagated from the PPP, MLPPP or Cisco HDLC will be squelched in favor of the higher priority Ethernet SAP failure. All legacy fault conditions, including admin port down, will in turn be dismissed for the duration of the Ethernet fault and will not be rediscovered until the expiration of the recovery-timer. This configurable timer value is the amount of time the process waits to allow the legacy connections to recover and establish following the clearing of the Ethernet fault. If the timer value is too short then false positive propagation will occur from the legacy side to the Ethernet connection. If the timer value is too long then secondary legacy faults will not be propagated to the associated Ethernet SAP for an extended period of time, delaying the proper state on the layer three device connected to the Ethernet SAP. Any packets arriving on the Ethernet SAP will be dropped until the legacy connection has recovered. As soon as the legacy connection recovers forwarding across the lpipe will occur regardless of the amount of time remaining for the recovery timer. Operators are required to adjust this timer value to their specific network requirements. If the timer adjustment is made while the service is active, the new timer will replace the old value and the new value will start counting down when called.

If the **eth-legacy-fault-notification** command is disabled from an active lpipe service then any previously reported fault will be cleared and the recovery-timer will be started. If the **eth-legacy-fault-notification** command is added to an active lpipe service, the process will check for outstanding faults and take the appropriate action.

Cisco HDLC behavior must be modified in order to better align with the fault inter-working function. In order to enable the **eth-legacy-fault-notification**, keepalives must be enabled. The following describes the new behavior for the Cisco HDLC port:

- Operationally up if it is receiving keepalives and has physical link (same behavior in either case)
- Operationally up if keepalives are disabled locally and has physical link (irrelevant for this feature because keepalives must be enabled). This is included for completeness.
- Operationally down when no keepalives are received and keepalives are locally enabled (same behavior in either case)
- Operationally down when there is no physical port (same behavior in either case)
- Operationally Down if it is part of a SAP but there no ce-address and has physical link (altered behavior)
- Operationally Down if it is part of a SAP but the SAP is shutdown and has physical link (altered behavior)
- Operationally Down if it is part of a SAP and the service is shutdown and has physical link (altered behavior)

The show service command has been expanded to include the basic Ethernet Legacy Fault Notification information and the specific SAP configuration.

The “Eth Legacy Fault Notification” section displays the configured recovery-timer value and whether the **eth-legacy-fault-notification** is active “**Admin State: inService**” (no shutdown) or inactive “**Admin State: outOfService**” (shutdown).

The “Ipipe SAP Configuration Information” displays the current Ethernet fault propagated to the associated legacy connection state; “**Legacy Fault Notify**”: False indicates no fault is currently being propagated and **True** indicates fault is currently being propagated. The “**Recvry Timer Rem**” is used to show the amount of time remaining before the recovery timer expires. A time in seconds will only be displayed for this parameter if an Ethernet fault has cleared and the recovery timer is currently counting down to 0.0 seconds.

A number of examples have been included using the service configuration below to demonstrate the various conditions. Many of the display commands have been trimmed in an effort to present feature relevant information.

```
configure service ipipe 201
    description "IPIPE_PPP"
    service-mtu 1514
    eth-legacy-fault-notification
        recovery-timer 300
        no shutdown
    exit
    ce-address-discovery ipv6 keep
    service-name "XYZ Ipipe 201"
    sap 1/1/4:21 create
        description "Default sap description for service id 201"
        eth-cfm
            mep 22 domain 1 association 45 direction down
            fault-propagation-enable use-if-tlv
            ccm-enable
            no shutdown
        exit
    exit
    exit
    sap 2/2/1.1.2.1 create
        description "Default sap description for service id 201"
    exit
    no shutdown
```

Service fully operational with no faults.

```
show service id 201 all
=====
Service Detailed Information
=====
Service Id      : 201                Vpn Id          : 201
Service Type    : Ipipe
Name            : XYZ Ipipe 201
Description     : IPIPE_PPP
```

```

Customer Id      : 1                      Creation Origin   : manual
Last Status Change: 01/07/2015 15:07:54
Last Mgmt Change : 01/07/2015 15:07:53
Admin State      : Up                      Oper State        : Up
MTU              : 1514
Vc Switching     : False
SAP Count        : 2                      SDP Bind Count     : 0
CE IPv4 Discovery : Enabled                Keep address       : Yes
CE IPv6 Discovery : Enabled                Stack Cap Sig      : Disabled

```

Eth Legacy Fault Notification

```

-----
Recovery Timer    : 30.0 secs              Admin State       : inService
-----

```

ETH-CFM service specifics

```

-----
Tunnel Faults     : ignore
-----

```

Service Destination Points (SDPs)

```

-----
No Matching Entries
-----

```

Service Access Points

SAP 1/1/4:21

```

-----
Service Id        : 201
SAP               : 1/1/4:21              Encap              : q-tag
Description       : Default sap description for service id 201
Admin State       : Up                    Oper State         : Up
Flags             : None
Multi Svc Site    : None
Last Status Change: 01/07/2015 15:07:53
Last Mgmt Change  : 01/07/2015 15:07:52
Sub Type          : regular
Dot1Q Ethertype   : 0x8100                QinQ Ethertype     : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU         : 1518                  Oper MTU           : 1518
Ingr IP Fltr-Id   : n/a                  Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id  : n/a                  Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a                  Egr IPv6 Fltr-Id  : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint          : N/A
Q Frame-Based Acct : Disabled              Limit Unused BW    : Disabled
Agg Burst Limit   : default

Acct. Pol         : None                  Collect Stats      : Disabled

Application Profile: None
Transit Policy    : None

Oper Group        : (none)                Monitor Oper Grp   : (none)
Host Lockout Plcy : n/a
Ignore Oper Down  : Disabled
Lag Link Map Prof : (none)

```

```

-----
ETH-CFM SAP specifics
-----
Tunnel Faults      : n/a                      AIS                      : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : None
-----

Ipipe SAP Configuration Information
-----
Configured CE IPv4 : n/a                      Discovered CE IPv4: 32.32.32.1
SAP MAC Address    : fe:ed:01:01:00:04        Mac Refresh Inter*: 14400
-----

Ipipe SAP IPv4 ARP Entry Info
-----
32.32.32.1          fe:4e:01:01:00:03 dynamic
-----

Ipipe SAP IPv6 Neighbor Entry Info
-----
fe80::fc2e:ffff:fe00:0    fe:4e:01:01:00:03 dynamic
3ffe::2020:2001           fe:4e:01:01:00:03 dynamic
. . . snip . . .
-----

Eth-Cfm MEP Configuration Information
-----
Md-index      : 1                      Direction      : Down
Ma-index      : 45                     Admin          : Enabled
MepId         : 22                     CCM-Enable     : Enabled
IfIndex       : 35782656               PrimaryVid     : 21
Description    : (Not Specified)
FngAlarmTime  : 0                      FngResetTime   : 0
FngState      : fngReset                ControlMep     : False
LowestDefectPri : macRemErrXcon         HighestDefect   : none
Defect Flags   : None
Mac Address    : fe:ed:01:01:00:04      Collect LMM Stats : disabled
CcmLtmPriority : 7                      CcmPaddingSize  : 0 octets
CcmTx         : 471                     CcmSequenceErr  : 0
CcmIgnoreTLVs : (Not Specified)
Fault Propagation : useIfStatusTLV      FacilityFault    : n/a
MA-CcmInterval : 1                      MA-CcmHoldTime  : 0ms
MA-Primary-Vid : Disabled
Eth-1Dm Threshold : 3(sec)              MD-Level        : 1
Eth-Ais       : Disabled
Eth-Ais Tx defCCM : allDef
Eth-Tst       : Disabled
Eth-CSF       : Disabled

Redundancy:
  MC-LAG State : n/a
  LbRxReply    : 0                      LbRxBadOrder    : 0
  LbRxBadMsdu  : 0                      LbTxReply       : 0
  LbSequence   : 1                      LbNextSequence  : 1
  LbRxUnexplained : 0
* indicates that the corresponding row element may have been truncated.

```

```

-----
SAP 2/2/1.1.2.1
-----
Service Id      : 201
SAP             : 2/2/1.1.2.1      Encap             : ipcp
Description     : Default sap description for service id 201
Admin State     : Up               Oper State       : Up
Flags          : None
Multi Svc Site  : None
Last Status Change : 01/07/2015 15:08:03
Last Mgmt Change  : 01/07/2015 15:07:54
Sub Type        : regular
Split Horizon Group: (Not Specified)

Admin MTU       : 1600              Oper MTU        : 1600
Ingr IP Fltr-Id : n/a              Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a             Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a            Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint        : N/A
Limit Unused BW : Disabled

Agg Burst Limit : default

Acct. Pol       : None              Collect Stats    : Disabled

Application Profile: None
Transit Policy   : None

Oper Group       : (none)           Monitor Oper Grp : (none)
Host Lockout Plcy : n/a
Ignore Oper Down : Enabled
Lag Link Map Prof : (none)
-----
Ipipe SAP Configuration Information
-----
Configured CE IPv4 : n/a           Discovered CE IPv4: 0.0.0.0
Legacy Fault Notify: False         Recvry Timer Rem  : 0.0 secs
-----
Ipipe SAP IPv4 ARP Entry Info
-----
No Ipipe SAP IPv4 ARP entries
-----
Ipipe SAP IPv6 Neighbor Entry Info
-----
fe80::13:9295:9ba:5e2             dynamic

. . . snip . . .

show port 2/2/1.1.2.1
=====
TDM DS0 Chan Group
=====
Description      : DS0GRP
Interface        : 2/2/1.1.2.1

```



```

TimeSlots      : 2-32
Speed          : 64
Admin Status   : up
BER SF Link Down : disabled
Last State Change : 01/07/2015 15:08:09
Configured Address : fe:ee:02:02:00:01
Hardware Address  : fe:ee:02:02:00:01
CRC            : 16
Oper Status    : up
Chan-Grp IfIndex : 608206967

```

```

Configured mode : access
Admin MTU       : 1600
Scramble        : false
Physical Link    : yes
Idle Cycle Flags : flags
Payload Fill Type : n/a
Signal Fill Type : n/a
Ing. Pool % Rate : 100
Egr. Sched. Pol : N/A
Encap Type      : ipcp
Oper MTU        : 1600
Bundle Number   : none
Load-balance-algo : Default
Payload Pattern : N/A
Signal Pattern   : N/A
Egr. Pool % Rate : 100

```

Traffic Statistics

	Input	Output
Octets	117200	246356
Packets	983	1004
Errors	0	0

Port Statistics

	Input	Output
Packets	983	1004
Discards	0	0
Unknown Proto Discards	0	

show port 2/2/1.1.2.1 ppp

PPP Protocols for 2/2/1.1.2.1

Protocol	State	Last Change	Restart Count	Last Cleared
lcp	opened	01/07/2015 15:08:08	1	01/07/2015 15:07:22
ipcp	opened	01/07/2015 15:08:08	1	01/07/2015 15:07:22
mplscp	initial	11/30/2014 09:20:08	0	01/07/2015 15:07:22
bcp	initial	11/30/2014 09:20:08	0	01/07/2015 15:07:22
osicp	initial	11/30/2014 09:20:08	0	01/07/2015 15:07:22
ipv6cp	opened	01/07/2015 15:08:20	1	01/07/2015 15:07:22

PPP Statistics

```

Local Mac address : fe:ee:02:02:00:01 Remote Mac address :
Local Magic Number : 0x7cda9060 Remote Magic Number: 0x23b8f81
Local IPv4 address : 32.32.32.1 Remote IPv4 address: 32.32.32.2
Local IPv6 address : fe80::fc2e:ffff:fe00:0

```

Remote IPv6 address: fe80::13:9295:9ba:5e2

Line Monitor Method: keepalive

Keepalive statistics

```
Request interval : 10          Threshold exceeded : 0
Drop Count       : 3          In packets           : 48
Time to link drop : 00h00m30s Out packets          : 48
Last cleared time : 01/07/2015 15:07:22
```

PPP Header Compression

ACFC : Disabled PFC : Disabled

show service sap-using

Service Access Points

PortId	SvcId	Ing. QoS	Ing. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
1/1/4:21	201	1	none	1	none	Up	Up
2/2/1.1.2.1	201	1	none	1	none	Up	Up

Number of SAPs : 8

The same service is used to demonstrate an Ethernet SAP failure condition propagating fault to the associated PPP connection. In this case an ETH-CCM time out has occurred. Only the changes have been highlighted.

The log events below will be specific to the failure type and the protocols involved.

166 2015/01/07 15:18:07.26 UTC MINOR: ETH_CFM #2001 Base

"MEP 1/45/22 highest defect is now defRemoteCCM"

167 2015/01/07 15:18:07.31 UTC MINOR: PPP #2004 Base 2/2/1.ds0grp-1.2.1

"Port 2/2/1.ds0grp-1.2.1 ipcp left 'opened' state"

168 2015/01/07 15:18:07.31 UTC MINOR: PPP #2004 Base 2/2/1.ds0grp-1.2.1

"Port 2/2/1.ds0grp-1.2.1 ipv6cp left 'opened' state"

169 2015/01/07 15:18:07.30 UTC MINOR: PPP #2002 Base 2/2/1.ds0grp-1.2.1

"Port 2/2/1.ds0grp-1.2.1 lcp left 'opened' state"

170 2015/01/07 15:18:07.30 UTC WARNING: SNMP #2004 Base 2/2/1.ds0grp-1.2.1

"Interface 2/2/1.ds0grp-1.2.1 is not operational"

171 2015/01/07 15:18:07.30 UTC MAJOR: SVCMGR #2210 Base

"Processing of an access port state change event is finished and the status of all affected SAPs on port 2/2/1.1.2.1 has been updated."

show service id 201 all

Service Detailed Information

Service Id : 201 Vpn Id : 201

```

Service Type      : Ipipe
Name              : XYZ Ipipe 201
Description       : IPIPE_PPP
Customer Id      : 1                      Creation Origin   : manual
Last Status Change: 01/07/2015 15:07:54
Last Mgmt Change  : 01/07/2015 15:07:53
Admin State      : Up                      Oper State        : Up
MTU              : 1514
Vc Switching     : False
SAP Count        : 2                      SDP Bind Count    : 0
CE IPv4 Discovery : Enabled                Keep address      : Yes
CE IPv6 Discovery : Enabled                Stack Cap Sig     : Disabled

Eth Legacy Fault Notification
-----
Recovery Timer    : 30.0 secs              Admin State       : inService
-----

ETH-CFM service specifics
-----
Tunnel Faults     : ignore
-----

Service Destination Points (SDPs)
-----
No Matching Entries
-----

Service Access Points
-----

-----
SAP 1/1/4:21
-----
Service Id        : 201
SAP               : 1/1/4:21              Encap             : q-tag
Description       : Default sap description for service id 201
Admin State      : Up                      Oper State        : Up
Flags            : OamDownMEPFault
Multi Svc Site   : None
Last Status Change: 01/07/2015 15:07:53
Last Mgmt Change  : 01/07/2015 15:07:52
Sub Type         : regular
Dot1Q Ethertype  : 0x8100                 QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU         : 1518                    Oper MTU          : 1518
Ingr IP Fltr-Id   : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id  : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a                     Egr IPv6 Fltr-Id  : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint          : N/A
Q Frame-Based Acct : Disabled                Limit Unused BW   : Disabled
Agg Burst Limit   : default

Acct. Pol         : None                      Collect Stats     : Disabled

Application Profile: None

```

```

Transit Policy      : None

Oper Group          : (none)                Monitor Oper Grp : (none)
Host Lockout Plcy   : n/a
Ignore Oper Down    : Disabled
Lag Link Map Prof   : (none)

-----
ETH-CFM SAP specifics
-----
Tunnel Faults       : n/a                    AIS                : Disabled
MC Prop-Hold-Timer  : n/a
Squelch Levels      : None

-----
Ipipe SAP Configuration Information
-----
Configured CE IPv4   : n/a                    Discovered CE IPv4: 32.32.32.1
SAP MAC Address      : fe:ed:01:01:00:04      Mac Refresh Inter*: 14400

-----
Ipipe SAP IPv4 ARP Entry Info
-----
32.32.32.1           fe:4e:01:01:00:03 dynamic

-----
Ipipe SAP IPv6 Neighbor Entry Info
-----
fe80::fc2e:ffff:fe00:0   fe:4e:01:01:00:03 dynamic
3ffe::2020:2001          fe:4e:01:01:00:03 dynamic

. . . snip . . .

-----
Eth-Cfm MEP Configuration Information
-----
Md-index            : 1                      Direction           : Down
Ma-index            : 45                     Admin               : Enabled
MepId               : 22                     CCM-Enable          : Enabled
IfIndex             : 35782656               PrimaryVid           : 21
Description          : (Not Specified)
FngAlarmTime        : 0                      FngResetTime        : 0
FngState             : fngDefectReported      ControlMep           : False
LowestDefectPri      : macRemErrXcon           HighestDefect         : defRemoteCCM
Defect Flags         : bDefRemoteCCM
Mac Address          : fe:ed:01:01:00:04      Collect LMM Stats    : disabled
CcmLtmPriority        : 7                     CcmPaddingSize       : 0 octets
CcmTx                : 650                    CcmSequenceErr       : 0
CcmIgnoreTLVs        : (Not Specified)
Fault Propagation     : useIfStatusTLV         FacilityFault         : n/a
MA-CcmInterval       : 1                     MA-CcmHoldTime       : 0ms
MA-Primary-Vid       : Disabled
Eth-ldm Threshold    : 3(sec)                 MD-Level              : 1
Eth-Ais              : Disabled
Eth-Ais Tx defCCM    : allDef
Eth-Tst              : Disabled
Eth-CSF              : Disabled

Redundancy:
  MC-LAG State        : n/a

```

```

LbRxReply          : 0
LbRxBadMsdu        : 0
LbSequence         : 1
LtRxUnexplained    : 0
LbRxBadOrder       : 0
LbTxReply          : 0
LbNextSequence     : 1
* indicates that the corresponding row element may have been truncated.

-----
SAP 2/2/1.1.2.1
-----
Service Id         : 201
SAP                : 2/2/1.1.2.1
Description        : Default sap description for service id 201
Admin State        : Up
Flags              : PortOperDown
Multi Svc Site     : None
Last Status Change : 01/07/2015 15:08:03
Last Mgmt Change   : 01/07/2015 15:07:54
Sub Type           : regular
Split Horizon Group: (Not Specified)

Admin MTU          : 1600
Ingr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a
Oper MTU           : 1600
Egr IP Fltr-Id     : n/a
Egr Mac Fltr-Id    : n/a
Egr IPv6 Fltr-Id   : n/a
qinq-pbit-marking  : both
Egr Agg Rate Limit: max

Endpoint           : N/A
Limit Unused BW    : Disabled

Agg Burst Limit    : default

Acct. Pol          : None
Collect Stats      : Disabled

Application Profile: None
Transit Policy     : None

Oper Group         : (none)
Host Lockout Plcy  : n/a
Ignore Oper Down   : Enabled
Lag Link Map Prof  : (none)
Monitor Oper Grp   : (none)

-----
Ipipe SAP Configuration Information
-----
Configured CE IPv4 : n/a
Legacy Fault Notify: True
Discovered CE IPv4: 0.0.0.0
Recvry Timer Rem   : 0.0 secs

-----
Ipipe SAP IPv4 ARP Entry Info
-----
No Ipipe SAP IPv4 ARP entries

-----
Ipipe SAP IPv6 Neighbor Entry Info
-----
No Ipipe SAP IPv6 Neighbor entries

. . . snip . . .

show port 2/2/1.1.2.1
=====

```

```

TDM DS0 Chan Group
=====
Description      : DS0GRP
Interface       : 2/2/1.1.2.1
TimeSlots      : 2-32
Speed          : 64
Admin Status    : up
BER SF Link Down : disabled
Last State Change : 01/07/2015 15:18:07
Configured Address : fe:ee:02:02:00:01
Hardware Address : fe:ee:02:02:00:01

Configured mode  : access
Admin MTU       : 1600
Scramble        : false
Physical Link    : yes
Idle Cycle Flags : flags
Payload Fill Type : n/a
Signal Fill Type : n/a
Ing. Pool % Rate : 100
Egr. Sched. Pol : N/A

Encap Type      : ipcp
Oper MTU       : 1600
Bundle Number   : none
Load-balance-algo : Default
Payload Pattern : N/A
Signal Pattern  : N/A
Egr. Pool % Rate : 100
Chan-Grp IfIndex : 608206967

=====

Traffic Statistics
=====
                        Input          Output
-----
Octets                117764         247052
Packets                1025           1034
Errors                  0              0

=====

Port Statistics
=====
                        Input          Output
-----
Packets                1025           1034
Discards                0              0
Unknown Proto Discards  0

=====
*A:Dut-B# show port 2/2/1.1.2.1 ppp

=====

PPP Protocols for 2/2/1.1.2.1
=====
Protocol  State      Last Change      Restart Count    Last Cleared
-----
lcp       initial    01/07/2015 15:18:07      1      01/07/2015 15:07:22
ipcp      initial    01/07/2015 15:18:07      1      01/07/2015 15:07:22
mplscp    initial    11/30/2014 09:20:08      0      01/07/2015 15:07:22
bcp       initial    11/30/2014 09:20:08      0      01/07/2015 15:07:22
osicp     initial    11/30/2014 09:20:08      0      01/07/2015 15:07:22
ipv6cp    initial    01/07/2015 15:18:07      1      01/07/2015 15:07:22

=====

PPP Statistics
=====

```

```
Local Mac address : fe:ee:02:02:00:01 Remote Mac address :
Local Magic Number : 0x0 Remote Magic Number: 0x0
Local IPv4 address : 0.0.0.0 Remote IPv4 address: 0.0.0.0
Local IPv6 address : ::
Remote IPv6 address: ::
```

Line Monitor Method: keepalive

Keepalive statistics

```
Request interval : 10 Threshold exceeded : 0
Drop Count : 3 In packets : 61
Time to link drop : 00h00m30s Out packets : 61
Last cleared time : 01/07/2015 15:07:22
```

PPP Header Compression

```
ACFC : Disabled PFC : Disabled
```

When the Ethernet fault condition clears a transitional state occurs.

```
172 2015/01/07 15:34:33.32 UTC MINOR: ETH_CFM #2001 Base
"MEP 1/45/22 highest defect is now none"
```

show service id 201 all

Service Detailed Information

```
Service Id : 201 Vpn Id : 201
Service Type : Ipipe
Name : XYZ Ipipe 201
Description : IPIPE_PPP
Customer Id : 1 Creation Origin : manual
Last Status Change: 01/07/2015 15:07:54
Last Mgmt Change : 01/07/2015 15:07:53
Admin State : Up Oper State : Up
MTU : 1514
Vc Switching : False
SAP Count : 2 SDP Bind Count : 0
CE IPv4 Discovery : Enabled Keep address : Yes
CE IPv6 Discovery : Enabled Stack Cap Sig : Disabled
```

Eth Legacy Fault Notification

```
Recovery Timer : 30.0 secs Admin State : inService
```

ETH-CFM service specifics

```
Tunnel Faults : ignore
```

Service Destination Points (SDPs)

No Matching Entries

Service Access Points

```

-----
SAP 1/1/4:21
-----
Service Id      : 201
SAP             : 1/1/4:21          Encap             : q-tag
Description     : Default sap description for service id 201
Admin State     : Up                Oper State       : Up
Flags           : None
Multi Svc Site  : None
Last Status Change : 01/07/2015 15:07:53
Last Mgmt Change  : 01/07/2015 15:07:52
Sub Type        : regular
Dot1Q Ethertype : 0x8100            QinQ Ethertype   : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU       : 1518              Oper MTU         : 1518
Ingr IP Fltr-Id : n/a              Egr IP Fltr-Id   : n/a
Ingr Mac Fltr-Id : n/a             Egr Mac Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a            Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint        : N/A
Q Frame-Based Acct : Disabled       Limit Unused BW  : Disabled
Agg Burst Limit  : default

Acct. Pol       : None              Collect Stats    : Disabled

Application Profile: None
Transit Policy   : None

Oper Group       : (none)           Monitor Oper Grp : (none)
Host Lockout Plcy : n/a
Ignore Oper Down  : Disabled
Lag Link Map Prof : (none)

-----
ETH-CFM SAP specifics
-----
Tunnel Faults    : n/a              AIS              : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels   : None

-----
Ipipe SAP Configuration Information
-----
Configured CE IPv4 : n/a            Discovered CE IPv4: 32.32.32.1
SAP MAC Address    : fe:ed:01:01:00:04 Mac Refresh Inter*: 14400

-----
Ipipe SAP IPv4 ARP Entry Info
-----
32.32.32.1       fe:4e:01:01:00:03 dynamic

-----
Ipipe SAP IPv6 Neighbor Entry Info
-----
fe80::fc2e:ffff:fe00:0 fe:4e:01:01:00:03 dynamic
3ffe::2020:2001      fe:4e:01:01:00:03 dynamic

```


. . . snip . . .

Eth-Cfm MEP Configuration Information

Md-index	: 1	Direction	: Down
Ma-index	: 45	Admin	: Enabled
MepId	: 22	CCM-Enable	: Enabled
IfIndex	: 35782656	PrimaryVid	: 21
Description	: (Not Specified)		
FngAlarmTime	: 0	FngResetTime	: 0
FngState	: fngReset	ControlMep	: False
LowestDefectPri	: macRemErrXcon	HighestDefect	: none
Defect Flags	: None		
Mac Address	: fe:ed:01:01:00:04	Collect LMM Stats	: disabled
CcmLtmPriority	: 7	CcmPaddingSize	: 0 octets
CcmTx	: 1603	CcmSequenceErr	: 0
CcmIgnoreTLVs	: (Not Specified)		
Fault Propagation	: useIfStatusTLV	FacilityFault	: n/a
MA-CcmInterval	: 1	MA-CcmHoldTime	: 0ms
MA-Primary-Vid	: Disabled		
Eth-IDm Threshold	: 3(sec)	MD-Level	: 1
Eth-Ais	: Disabled		
Eth-Ais Tx defCCM	: allDef		
Eth-Tst	: Disabled		
Eth-CSF	: Disabled		

Redundancy:

MC-LAG State	: n/a		
LbRxReply	: 0	LbRxBadOrder	: 0
LbRxBadMsdu	: 0	LbTxReply	: 0
LbSequence	: 1	LbNextSequence	: 1
LtRxUnexplained	: 0		

* indicates that the corresponding row element may have been truncated.

SAP 2/2/1.1.2.1

Service Id	: 201		
SAP	: 2/2/1.1.2.1	Encap	: ipcp
Description	: Default sap description for service id 201		
Admin State	: Up	Oper State	: Up
Flags	: PortOperDown		
Multi Svc Site	: None		
Last Status Change	: 01/07/2015 15:08:03		
Last Mgmt Change	: 01/07/2015 15:07:54		
Sub Type	: regular		
Split Horizon Group	: (Not Specified)		
Admin MTU	: 1600	Oper MTU	: 1600
Ingr IP Fltr-Id	: n/a	Egr IP Fltr-Id	: n/a
Ingr Mac Fltr-Id	: n/a	Egr Mac Fltr-Id	: n/a
Ingr IPv6 Fltr-Id	: n/a	Egr IPv6 Fltr-Id	: n/a
		qinq-pbit-marking	: both
		Egr Agg Rate Limit	: max
Endpoint	: N/A		
Agg Burst Limit	: default	Limit Unused BW	: Disabled

Acct. Pol : None Collect Stats : Disabled

Application Profile: None
Transit Policy : None

Oper Group : (none) Monitor Oper Grp : (none)
Host Lockout Plcy : n/a
Ignore Oper Down : Enabled
Lag Link Map Prof : (none)

Ipipe SAP Configuration Information

Configured CE IPv4 : n/a Discovered CE IPv4: 0.0.0.0
Legacy Fault Notify: False Recvry Timer Rem : 28.8 secs

Ipipe SAP IPv4 ARP Entry Info

No Ipipe SAP IPv4 ARP entries

Ipipe SAP IPv6 Neighbor Entry Info

No Ipipe SAP IPv6 Neighbor entries

. . . snip . . .

show port 2/2/1.1.2.1

=====

TDM DS0 Chan Group

=====

Description	: DS0GRP		
Interface	: 2/2/1.1.2.1		
TimeSlots	: 2-32		
Speed	: 64	CRC	: 16
Admin Status	: up	Oper Status	: down
BER SF Link Down	: disabled		
Last State Change	: 01/07/2015 15:18:07	Chan-Grp IfIndex	: 608206967
Configured Address	: fe:ee:02:02:00:01		
Hardware Address	: fe:ee:02:02:00:01		

Configured mode	: access	Encap Type	: ipcp
Admin MTU	: 1600	Oper MTU	: 1600
Scramble	: false		
Physical Link	: yes	Bundle Number	: none
Idle Cycle Flags	: flags	Load-balance-algo	: Default
Payload Fill Type	: n/a	Payload Pattern	: N/A
Signal Fill Type	: n/a	Signal Pattern	: N/A
Ing. Pool % Rate	: 100	Egr. Pool % Rate	: 100
Egr. Sched. Pol	: N/A		

=====

=====

Traffic Statistics

	Input	Output
-----	-----	-----
Octets	119518	247124

```

Packets                                1123                                1036
Errors                                0                                0

=====
Port Statistics
=====
                                Input                                Output
-----
Packets                                1123                                1036
Discards                                0                                0
Unknown Proto Discards                0
=====

show port 2/2/1.1.2.1 ppp
=====
PPP Protocols for 2/2/1.1.2.1
=====
Protocol  State          Last Change          Restart Count    Last Cleared
-----
lcp       request sent  01/07/2015 15:34:33      1    01/07/2015 15:07:22
ipcp      initial      01/07/2015 15:18:07      1    01/07/2015 15:07:22
mplscp    initial      11/30/2014 09:20:08      0    01/07/2015 15:07:22
bcp       initial      11/30/2014 09:20:08      0    01/07/2015 15:07:22
osicp     initial      11/30/2014 09:20:08      0    01/07/2015 15:07:22
ipv6cp    initial      01/07/2015 15:18:07      1    01/07/2015 15:07:22
=====

=====
PPP Statistics
=====
Local Mac address  : fe:ee:02:02:00:01  Remote Mac address :
Local Magic Number : 0x0              Remote Magic Number: 0x0
Local IPv4 address : 32.32.32.1        Remote IPv4 address: 0.0.0.0
Local IPv6 address : fe80::fc2e:ffff:fe00:0
Remote IPv6 address: ::

Line Monitor Method: keepalive

Keepalive statistics

Request interval   : 10              Threshold exceeded : 0
Drop Count         : 3                In packets         : 61
Time to link drop  : 00h00m30s        Out packets        : 61
Last cleared time  : 01/07/2015 15:07:22

PPP Header Compression
ACFC                : Disabled      PFC                : Disabled
=====

```

An example of the legacy fault propagation to the associated Ethernet SAP and the remote peer using the ETH-CFM fault propagation, assuming no Ethernet Fault is taking precedence.

```

173 2015/01/07 15:35:03.31 UTC MINOR: SVCMMGR #2203 Base
"Status of SAP 2/2/
1.1.2.1 in service 201 (customer 1) changed to admin=up oper=down flags=PortOperDown
"

```

```

show service id 201 all
=====
Service Detailed Information
=====
Service Id      : 201                Vpn Id          : 201
Service Type    : Ipipe
Name            : XYZ Ipipe 201
Description     : IPIPE_PPP
Customer Id     : 1                  Creation Origin  : manual
Last Status Change: 01/07/2015 15:07:54
Last Mgmt Change : 01/07/2015 15:07:53
Admin State     : Up                 Oper State      : Up
MTU             : 1514
Vc Switching    : False
SAP Count       : 2                  SDP Bind Count  : 0
CE IPv4 Discovery : Enabled           Keep address    : Yes
CE IPv6 Discovery : Enabled           Stack Cap Sig   : Disabled

Eth Legacy Fault Notification
-----
Recovery Timer   : 30.0 secs          Admin State     : inService

-----
ETH-CFM service specifics
-----
Tunnel Faults    : ignore

-----
Service Destination Points(SDPs)
-----
No Matching Entries

-----
Service Access Points
-----

-----
SAP 1/1/4:21
-----
Service Id      : 201
SAP             : 1/1/4:21           Encap           : q-tag
Description     : Default sap description for service id 201
Admin State     : Up                 Oper State      : Up
Flags           : None
Multi Svc Site  : None
Last Status Change: 01/07/2015 15:07:53
Last Mgmt Change : 01/07/2015 15:07:52
Sub Type        : regular
Dot1Q Ethertype : 0x8100             QinQ Ethertype  : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU       : 1518                Oper MTU        : 1518
Ingr IP Fltr-Id : n/a                 Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a                 Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a                 Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit: max

Endpoint        : N/A
Q Frame-Based Acct : Disabled          Limit Unused BW  : Disabled
Agg Burst Limit  : default

```

Acct. Pol : None Collect Stats : Disabled

Application Profile: None

Transit Policy : None

Oper Group : (none)

Monitor Oper Grp : (none)

Host Lockout Plcy : n/a

Ignore Oper Down : Disabled

Lag Link Map Prof : (none)

ETH-CFM SAP specifics

Tunnel Faults : n/a

AIS : Disabled

MC Prop-Hold-Timer : n/a

Squelch Levels : None

Ipipe SAP Configuration Information

Configured CE IPv4 : n/a

Discovered CE IPv4: 32.32.32.1

SAP MAC Address : fe:ed:01:01:00:04

Mac Refresh Inter*: 14400

Ipipe SAP IPv4 ARP Entry Info

32.32.32.1

fe:4e:01:01:00:03 dynamic

Ipipe SAP IPv6 Neighbor Entry Info

fe80::fc2e:ffff:fe00:0

fe:4e:01:01:00:03 dynamic

3ffe::2020:2001

fe:4e:01:01:00:03 dynamic

. . . snip . . .

Eth-Cfm MEP Configuration Information

Md-index : 1

Direction : Down

Ma-index : 45

Admin : Enabled

MepId : 22

CCM-Enable : Enabled

IfIndex : 35782656

PrimaryVid : 21

Description : (Not Specified)

FngAlarmTime : 0

FngResetTime : 0

FngState : fngReset

ControlMep : False

LowestDefectPri : macRemErrXcon

HighestDefect : none

Defect Flags : bDefRDICCM

Mac Address : fe:ed:01:01:00:04

Collect LMM Stats : disabled

CcmLtmPriority : 7

CcmPaddingSize : 0 octets

CcmTx : 1690

CcmSequenceErr : 0

CcmIgnoreTLVs : (Not Specified)

Fault Propagation : useIfStatusTLV

FacilityFault : n/a

MA-CcmInterval : 1

MA-CcmHoldTime : 0ms

MA-Primary-Vid : Disabled

Eth-lDm Threshold : 3(sec)

MD-Level : 1

Eth-Ais : Disabled

Eth-Ais Tx defCCM : allDef

Eth-Tst : Disabled

Eth-CSF : Disabled

Redundancy:

MC-LAG State : n/a

LbRxReply : 0

LbRxBadOrder : 0

LbRxBadMsdu : 0

LbTxReply : 0

LbSequence : 1

LbNextSequence : 1

LtRxUnexplained : 0

* indicates that the corresponding row element may have been truncated.

SAP 2/2/1.1.2.1

Service Id : 201

SAP : 2/2/1.1.2.1

Encap : ipcp

Description : Default sap description for service id 201

Admin State : Up

Oper State : Down

Flags : PortOperDown

Multi Svc Site : None

Last Status Change : 01/07/2015 15:35:03

Last Mgmt Change : 01/07/2015 15:07:54

Sub Type : regular

Split Horizon Group: (Not Specified)

Admin MTU : 1600

Oper MTU : 1600

Ingr IP Fltr-Id : n/a

Egr IP Fltr-Id : n/a

Ingr Mac Fltr-Id : n/a

Egr Mac Fltr-Id : n/a

Ingr IPv6 Fltr-Id : n/a

Egr IPv6 Fltr-Id : n/a

qinq-pbit-marking : both

Egr Agg Rate Limit: max

Endpoint : N/A

Limit Unused BW : Disabled

Agg Burst Limit : default

Acct. Pol : None

Collect Stats : Disabled

Application Profile: None

Transit Policy : None

Oper Group : (none)

Monitor Oper Grp : (none)

Host Lockout Plcy : n/a

Ignore Oper Down : Enabled

Lag Link Map Prof : (none)

Ipipe SAP Configuration Information

Configured CE IPv4 : n/a

Discovered CE IPv4: 0.0.0.0

Legacy Fault Notify: False

Recvry Timer Rem : 0.0 secs

Ipipe SAP IPv4 ARP Entry Info

No Ipipe SAP IPv4 ARP entries

Ipipe SAP IPv6 Neighbor Entry Info

No Ipipe SAP IPv6 Neighbor entries

. . . snip . . .

show port 2/2/1.1.2.1

TDM DS0 Chan Group

```

=====
Description      : DS0GRP
Interface        : 2/2/1.1.2.1
TimeSlots       : 2-32
Speed           : 64
Admin Status    : up
BER SF Link Down : disabled
Last State Change : 01/07/2015 15:18:07
Configured Address : fe:ee:02:02:00:01
Hardware Address  : fe:ee:02:02:00:01
  
```

```

Configured mode   : access
Admin MTU         : 1600
Scramble          : false
Physical Link     : yes
Idle Cycle Flags  : flags
Payload Fill Type : n/a
Signal Fill Type  : n/a
Ing. Pool % Rate  : 100
Egr. Sched. Pol   : N/A

Encap Type        : ipcp
Oper MTU          : 1600
Bundle Number     : none
Load-balance-algo : Default
Payload Pattern    : N/A
Signal Pattern     : N/A
Egr. Pool % Rate  : 100
  
```

Traffic Statistics

	Input	Output
Octets	119518	248132
Packets	1123	1064
Errors	0	0

Port Statistics

	Input	Output
Packets	1123	1064
Discards	0	0
Unknown Proto Discards	0	

show port 2/2/1.1.2.1 ppp

PPP Protocols for 2/2/1.1.2.1

Protocol	State	Last Change	Restart Count	Last Cleared
lcp	request sent	01/07/2015 15:36:05	1	01/07/2015 15:07:22
ipcp	initial	01/07/2015 15:18:07	1	01/07/2015 15:07:22
mplscp	initial	11/30/2014 09:20:08	0	01/07/2015 15:07:22
bcp	initial	11/30/2014 09:20:08	0	01/07/2015 15:07:22
osicp	initial	11/30/2014 09:20:08	0	01/07/2015 15:07:22
ipv6cp	initial	01/07/2015 15:18:07	1	01/07/2015 15:07:22

```

=====
PPP Statistics
=====
Local Mac address   : fe:ee:02:02:00:01  Remote Mac address :
Local Magic Number  : 0x0                Remote Magic Number: 0x0
Local IPv4 address  : 32.32.32.1         Remote IPv4 address: 0.0.0.0
Local IPv6 address  : fe80::fc2e:ffff:fe00:0
Remote IPv6 address: ::

Line Monitor Method: keepalive

Keepalive statistics

Request interval    : 10                Threshold exceeded : 0
Drop Count          : 3                 In packets          : 61
Time to link drop   : 00h00m30s         Out packets         : 61
Last cleared time   : 01/07/2015 15:07:22

PPP Header Compression
ACFC                : Disabled          PFC                : Disabled
=====

```

This feature is only supported for an Ipipe service that has a single legacy connection with an encaps-type PPP, MLPPP or Cisco-HDLC and an Ethernet SAP. No other combinations are supported. Deployments using APS cannot use this fault propagation functionality.

The propagation of fault is based on the interaction of a number of resources and software functions. This means that propagation and recovery will vary based on the type of failure, the scale of the failure, the legacy protocol, the system overhead at the time of the action, and other interactions.

Before maintenance operations are performed the operation should be aware of the operational state of the service and any fault propagation state. Admin legacy port state down conditions do not cause fault propagation, it is the operational port state that conveys fault. During a Major ISSU operation, legacy faults will be cleared and not propagated toward the Ethernet network. In order to prevent this clearing of faults, the operator may consider shutting down the Ethernet port or shutdown the ETH-CFM MEPs to cause a timeout upstream.



Note: The CLI commands for these functions can be found in the *L2 Services Guide*.

3.8.5.2 SAP/SDP-binding Failure (Including Pseudowire Status)

When a SAP/SDP-binding becomes faulty (oper-down, admin-down, or pseudowire status faulty), SMGR propagates the fault to OAM components on the mate SAP/SDP-binding.

3.8.5.3 Service Administratively Shutdown

When the service is administratively shutdown, SMGR propagates the fault to OAM components on both SAP/SDP-bindings.

3.8.5.4 Interaction with Pseudowire Redundancy

When the 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 only occurs when both pseudowires become faulty. Then the SMGR propagates the fault to CFM. Since there is no fault handling in the PIPE service, any CFM fault detected on a SDP-binding is not used in the pseudowire redundancy's algorithm to choose the most suitable SDP-binding to transmit on.

3.8.6 VPLS Service

For VPLS services, on down MEPs are supported for fault propagation.

3.8.6.1 CFM Detected Fault

When a MEP detects a fault and fault propagation is enabled for the MEP, CFM communicate the fault to the SMGR. The SMGR will mark the SAP/SDP-binding as oper-down. Note that oper-down is used here in VPLS instead of "oper-up but faulty" in the pipe services. CFM traffic can 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.

Note that as stated in [CFM Connectivity Fault Conditions](#), a fault is raised whenever a remote MEP is down (not all remote MEPs have to be down). When it is not desirable to trigger fault handling actions in some cases when a down MEP has multiple remote MEPs, operators can disable fault propagation for the MEP.

If the MEP is a down MEP, SMGR performs the fault handling actions for the affected service(s). Local actions done by the SMGR include (but are not limited to):

- Flushing MAC addresses learned on the faulty SAP/SDP-binding.
- Triggering transmission of MAC flush messages.
- Notifying MSTP/RSTP about topology change. If the VPLS instance is a management VPLS (mVPLS), all VPLS instances that are managed by the mVPLS inherits the MSTP/RSTP state change and react accordingly to it.
- If the service instance is a B-VPLS, and fault-propagation-bmac address(es) is/are configured for the SAP/SDP-binding, SMGR performs a lookup using the BMAC address(es) to find out which pipe services need to be notified, then propagates a fault to these services. There can be up to four remote BMAC addresses associated with an SAP/SDP-binding for the same B-VPLS.

3.8.6.2 SAP/SDP-Binding Failure (Including Pseudowire Status)

If the service instance is a B-VPLS, and an associated BMAC address is configured for the failed SAP/SDP-binding, the SMGR performs a lookup using the BMAC address to find out which pipe services will be notified and then propagate fault to these services.

Within the same B-VPLS service, all SAPs/SDP-bindings configured with the same fault propagation BMACs must be faulty or oper down for the fault to be propagated to the appropriate pipe services.

3.8.6.3 Service Down

When a VPLS service is down:

- If the service is not a B-VPLS service, the SMGR propagates the fault to OAM components on all SAP/SDP-bindings in the service.
- If the service is a B-VPLS service, the SMGR propagates the fault to OAM components on all SAP/SDP-bindings in the service as well as all pipe services that are associated with the B-VPLS instance.

3.8.6.4 Pseudowire Redundancy and Spanning Tree Protocol

A SAP or SDP binding that has a down MEP fault is made operationally down. This causes pseudowire redundancy or Spanning Tree Protocol (STP) to take the appropriate actions.

However, the reverse is not true. If the SAP or SDP binding is blocked by STP, or is not tx-active due to pseudowire redundancy, no fault is generated for this entity.

3.8.7 IES and VPRN Services

For IES and VPRN services, only down MEP is supported on Ethernet SAPs and spoke SDP bindings.

When a down MEP detects a fault and fault propagation is enabled for the MEP, CFM communicates the fault to the SMGR. The SMGR marks the SAP/SDP binding as operationally down. CFM traffic can still be transmitted to or received from the SAP/SDP-binding to ensure when the fault is cleared and the SAP will go back to normal operational state.

Because the SAP/SDP-binding goes down, it is not usable to upper applications. In this case, the IP interface on the SAP/SDP-binding go down. The prefix is withdrawn from routing updates to the remote PEs. The same applies to subscriber group interface SAPs on the 7450 ESS and 7750 SR.

When the IP interface is administratively shutdown, the SMGR notifies the down MEP and a CFM fault notification is generated to the CPE through interface status TLV or suspension of CCM based on local configuration.

3.8.8 Pseudowire Switching

When the node acts as a pseudowire switching node, meaning two pseudowires are stitched together at the node, the SMGR will not communicate pseudowire failures to CFM. Such features are expected to be communicated by pseudowire status messages, and CFM will run end-to-end on the head-end and tail-end of the stitched pseudowire for failure notification.

3.8.9 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.

3.8.10 802.3ah 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.

3.9 Service Assurance Agent (SAA)

Service Application Agent (SAA) is a tool that allows operators to configure a number of different tests that can be used to provide performance information like delay, jitter and loss for services or network segments. The test results are saved in SNMP tables or summarized XML files. These results can be collected and reported on using network management systems.

SAA uses the resources allocated to the various OAM processes. These processes are not dedicated to SAA but shared throughout the system. [Table 8](#) provides guidance on how these different OAM functions are logically grouped.

Table 8 SAA Test and Descriptions

Test	Description
Background	Tasks configured outside of the SAA hierarchy that consume OAM task resources. Specifically, these include SDP-Keep Alive, Static route cpe-check, filter redirect-policy, ping-test, and vrrp policy host-unreachable. These are critical tasks that ensure the network operation and may affect data forwarding or network convergence.
SAA Continuous	SAA tests configured as continuous (always scheduled).
SAA non-continuous	SAA tests that are not configured as continuous, hence scheduled outside of the SAA application. These tests require the oam saa test-name start command to initiate the test run.
Non-SAA (Directed)	Any task that does not include any configuration under SAA. These tests are SNMP or via the CLI that is used to troubleshoot or profile network condition. This would take the form “oam test-type” or ping/traceroute with the specific test parameters.

SAA test types are restricted to those that utilize a request response mechanism, single-ended tests. Dual-ended tests that initiate the test on one node but require the statistical gathering on the other node are not supported under SAA. As an example, Y.1731 defines two approaches for measuring frame delay and frame delay variation, single-ended and dual-ended. The single-ended approach is supported under SAA.

Post processing analysis of individual test runs can be used to determine the success or failure of the individual runs. The operator can set rising and lowering thresholds for delay, jitter, and loss. Exceeding the threshold will cause the test to have a failed result. A trap can be generated when the test fails. The operator is also able to configure a probe failure threshold and trap when these thresholds are exceeded.

Each supported test type has configuration properties specific to that test. Not all options, intervals, and parameters are available for all tests. Some configuration parameters, such as the sub second probe interval require specific hardware platforms.

The ETH-CFM SAA and **icmp-ping** (without the **rapid** parameter) test types may be configured as continuous, meaning automatically re-scheduled. When a test is executing, the probe spacing is based on the interval parameter, assuming there are no lost packets. When the **continuous** option is configured for supported tests, the probe spacing between the runs does not align with the configured **interval**. A number of closure and rescheduling functions occur that affect the probe spacing between runs.

Trace type tests apply the timeout to each individual packet, which may affect spacing. This is required because packet timeout may be required to move from one probe to the next probe. For tests that do not require this type of behavior, typically ping and ETH-CFM PM functions, the probes will be sent at the specified probe interval and the timeout is only applied at the end of the test if any probe has been lost during the run. When the timeout is applied at the end of the run, the test is considered complete when either all response have been received or the timeout expires at the end of the test run. For tests marked as continuous (always scheduled), the spacing between the runs may be delayed by the timeout value when a packet is lost. The test run is complete when all probes have either been received back or the timeout value has expired.

In order to preserve system resources, specifically memory, the operator should only store summarized history results. By default, summary results are stored for tests configured with sub second probe intervals, or a probe count above 100 or is written to a file. By default, per probe information will be stored for test configured with an interval of one second or above counters, and probe counts of 100 or less and is not written to a file. The operator may choose to override these defaults using the **probe-history {keep | drop | auto}** option. The **auto** option sets the defaults above. The other options override the default retention schemes based on the operator requirements, per probe retention **keep** or summary only information **drop**. The probe data can be viewed using the **show saa test** command. If the per probe information is retained, this probe data is available at the completion of the test run.

The summary data is updated throughout the test run. The overall memory system usage is available using the **show system memory-pools** command. The OAM entry represents the overall memory usage. This includes the history data stored for SAA tests. A **clear saa testname** option is available to release the memory and flush the test results.

The following example shows Y.1731 ETH-DMM packets to be sent from the local MEP 325, domain 12 and association 300 to destination MAC address d0:0d:1e:00:00:27. The tests will be scheduled as continuous and do not require an **oam saa test-name start** command to be issued by the operator. Each individual test run will contain 900 probes at 1 second intervals. This means each individual test run will be active for 15 minutes. If a packet is lost, the test will wait for the timeout (default 5s not shown) before closing one run and move to the next. If more than 10 probes are lost, the test will be marked as failed and a trap and log entry will be generated.

Test summary information and not per probe data is maintained for this test because the optional probe-history override is not configured. The summary information will be written to an XML file using the accounting-policy 1.

Example:

```
saa>test# info
-----
description "Two Way ETH-DDM To MEP 327 From MEP 325"
type
    eth-cfm-two-way-delay d0:0d:1e:00:00:27 mep 325 domain 12 association 300
    count 900 interval 1
exit
trap-gen
    probe-fail-enable
    probe-fail-threshold 10
exit
accounting-policy 1
continuous
no shutdown
```

SAA leverages the accounting record infrastructure. The sample configuration is included for completeness. For complete information on Accounting Policies consult the *7450 ESS, 7750 SR, and 7950 XRS System Management Guide*.

```
config>log# info
-----
file-id 1
    location cf3:
    rollover 60 retention 24
exit
accounting-policy 1
    description "SAA XML File"
    record saa
    collection-interval 15
    to file 1
```

```
no shutdown
exit
```

SAA launched tests will maintain two most recent completed and one in progress test. The output below is the summary data from the test above. Below, test run 18 and 19 have been completed and test run 20 is in progress. Once test run 20 is completed test run 18 data will be overwritten. It is important to ensure that the collection and accounting record process is configured in such a way to write the data to file before it is overwritten. Once the results are overwritten they are lost.

```
show saa "saa-dmm-1"
```

```
=====
SAA Test Information
=====
Test name           : saa-dmm-1
Owner name          : TiMOS CLI
Description          : Two Way ETH-DDM To MEP 327 From MEP 325
Accounting policy    : 1
Continuous          : Yes
Administrative status : Enabled
Test type           : eth-cfm-two-way-delay d0:0d:1e:00:00:27 mep
                    : 325 domain 12 association 300 count 900
                    : interval 1
Trap generation      : probe-fail-enable probe-fail-threshold 10
Probe History        : auto (drop)
Test runs since last clear : 3
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
           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: 18
Total number of attempts: 900
Number of requests that failed to be sent out: 0
Number of responses that were received: 900
=====
```


Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0

(in ms)	Min	Max	Average	Jitter
Outbound :	-29.3	-28.6	-28.9	0.000
Inbound :	28.7	29.3	29.0	0.000
Roundtrip :	0.069	0.077	0.073	0.000

Per test packet:

Test Run: 19
Total number of attempts: 900
Number of requests that failed to be sent out: 0
Number of responses that were received: 900
Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0

(in ms)	Min	Max	Average	Jitter
Outbound :	-29.9	-29.3	-29.6	0.000
Inbound :	29.3	30.0	29.7	0.001
Roundtrip :	0.069	0.080	0.073	0.001

Per test packet:

Test Run: 20
Total number of attempts: 181
Number of requests that failed to be sent out: 0
Number of responses that were received: 181
Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0

(in ms)	Min	Max	Average	Jitter
Outbound :	-30.0	-29.9	-30.0	0.001
Inbound :	30.0	30.1	30.0	0.000
Roundtrip :	0.069	0.075	0.072	0.001

Per test packet:

=====

Any data not written to file will be lost on a CPU switch over.

There are a number of **show** commands to help the operator monitor the test oam tool set.

show test-oam oam-config-summary —Provides information about the configured tests.

show test-oam oam-perf — Provides the transmit (launched form me) rate information and remotely launched test receive rate on the local network element.

clear test-oam oam-perf — Provides the ability to clear the test oam performance stats for a current view of the different rates in the **oam-perf** command above.

monitor test-oam oam-perf — Makes use of the **monitor** command to provide time sliced performance stats for test oam functions.


3.10 OAM Performance Monitoring (OAM-PM)

OAM Performance Monitoring (OAM-PM) provides an overall architecture for gathering and computing key performance indicators (KPI) using standard protocols and a robust collection model. The architecture is comprised of a number of foundational components.

1. Session: This is the overall collection of different tests, test parameters, measurement intervals, and mapping to configured storage models. It is the overall container that defines the attributes of the session.
2. Standard PM Packets: The protocols defined by various standards bodies that contains the necessary fields to collect statistical data for the performance attribute they represent. OAM-PM leverages single ended protocols. Single ended protocols follow a message response model, message sent by a launch point, response updated, and reflected by a responder.
3. Measurement Intervals (MI): Time based non-overlapping windows that captures all the results that are received in that window of time.
4. Data Structures: The unique counters and measurement results that represent the specific protocol.
5. Bin group: Ranges in micro seconds that counts the results that fit into the range.

The hierarchy of the architecture is captured in the [Figure 41](#). This diagram is only meant to draw the relationship between the components. It is not meant to depict all the detailed parameters required

Figure 41 OAM-PM Architecture Hierarchy



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 have its state shutdown. 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 sources and destinations configuration parameters are not validated by the test that makes uses of that information. Once the test is activated with a **no shutdown**, 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 or terminate the packets. If the session is a MEP-based Ethernet session and the source-based MEP does not exist, the transmit count for the test will be zero. If the source-based session is TWAMP Light, the OAM-PM transmit counter will increment but the receive counter will not.

OAM-PM is not a hitless operation. If a high availability event occurs, causing the backup CPM to become the newly active or when ISSU functions are performed. Tests in flight will not be completed, open files may not be closed, and test data not written to a properly closed XML file will be lost. 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](#)
- [Standard PM Packets](#)
- [Detectable Transmit Errors](#)
- [Measurement Intervals](#)
- [Data Structures and Storage](#)
- [Bin Groups](#)
- [Relating the Components](#)
- [IP Performance Monitoring](#)
- [Ethernet Performance Monitoring](#)

- [OAM-PM Event Monitoring](#)

3.10.1 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: Assigns the mantra of the test to either proactive (default) or on-demand. Individual test timing parameters will be influenced by this setting. A proactive session will start immediately following the **no shutdown** of the test. A proactive test will continue to execute until a manual shutdown stops the individual test. On-demand tests do not start immediately following the **no shutdown** command. The operator must start an on-demand test by using the command **oam>oam-pm>session>start** and specifying the applicable protocol. The operator can override the no test-duration default by configuring a fixed amount of time 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 and delete the tests when they are completed and all the results collected. This will free all system memory that has been reserved for storing the results. In the event of a high-availability event that causes the backup CPM to become the newly active, all on-demand tests will need to be restarted manually using the **oam>oam-pm>session>start** command for the specific protocol.

Test Family: The main branch of testing that will be addressed a specific technology. The available test parameters for the session will be based off the test family. The destination, source, and the priority are common to all tests under the session and defined separately from the individual test parameters.

Test Parameters: The parameters include individual tests with the associated parameters including start and stop times and the ability to activate and deactivate the individual test.

Measurement Interval: 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 may fail to be activated with a **no shutdown** command. If this situation occurs an error, it will indicate with “Invalid session parameters”.

3.10.2 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 protocol may be solely focused on measuring a single specific performance criteria or multiple. The protocols available to carry out the measurements will be 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, round trip, forward, and backward.

DMMv1 adds the following to the original DMM definition:

- Flag Field (1 bit – LSB) is defined as the Type (Proactive=1 | On-Demand=0)
- TestID TLV (32 bits) – 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 the Synthetic Loss Measurement (SLM) defined in Y.1731 by the ITU-T. This allows for the calculation of Frame Loss Ratio (FLR) and availability. The ITU-T also defines a frame loss measurement (LMM) approach that provides frame loss ratio (FLR) and raw transmit and receive frame counters in each direction and availability metrics.

IP Performance data uses the TWAMP test packet for gathering both delay and loss metrics. OAM-PM supports Appendix I of RFC 5357 (TWAMP Light). The SR OS supports the gathering of delay metrics Frame Delay (FD), InterFrame Delay Variation (IFDV), Frame Delay Range (FDR) and Mean Frame Delay (MFD) round trip, forward and backward.

A session can be configured with one test or multiple tests. Depending on sessions 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 within a session will share the common session parameters and 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. IP performance measurement uses a single TWAMP packet for both delay and synthetic loss.

Each test must be configured with a *TestID* as part of the test parameters. This 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 tuple *TestID*, Source MAC, and Destination MAC are used to maintain the uniqueness on the responder. This means the *TestID* has only local and not global significance. TWAMP test packets also require a *TestID* to be configured but do not carry this information in the PDU. However, it is required for uniform provisioning under the OAM-PM architecture. TWAMP uses a four tuple Source IP, Destination IP, Source UDP, and Destination UDP to maintain unique session indexes.

3.10.3 Detectable Transmit Errors

Each OAM-PM session test contains both an administrative and an operational state. The administrative state is linked to the **shutdown** or **no shutdown** configuration for the test. The operational state indicates if the test is actively sending or trying to send test frames. For all proactive session types, the administrative and operational states are linked. The test immediately starts and attempts to send test frames as soon as the **no shutdown** command is accepted. Sessions configured using the **session-type on-demand** command do not link the two states. The operational state is up when the test is actively generating or attempting to generate test frames. This does not occur at the time the **no shutdown** command is accepted. The **oam oam-pm ... start** command is required to commence the transmission of the test frames for on-demand tests. When the manually issued start command is accepted, the operational state changes to up. At the expiration of the test duration, or when the test is manually stopped using the **oam oam-pm ... stop** command, the operational state transitions to down.

A new field, "Detectable Tx Err", has been added to each Ethernet and IP OAM-PM test to indicate if there is a detectable transmit error that is preventing the test frames from being sent. This does not affect the two existing states. Detectable Tx Err information can be used to assist in troubleshooting, as it conveys information about a current detectable error state. No log events are created and no historical references are maintained when the error condition clears or changes.

The Detectable Tx Err condition is checked for each test which is operationally up and attempting to send frames. The audit function checks to see if a transmit error condition exists that could prevent the packet from being sent; for example, the source MEP is not fully configured, or the IP interface associated with the source is down. A raised detectable transmit error is cleared for the test if the audit process executes and finds no further detectable transmit error conditions. The interval of the audit function depends on a number of factors, such as the test family, the probe interval, and the number of active tests.

This is an ongoing maintenance function that executes while the test is operationally up. A maximum of one detectable transmit error can be presented to the operator. However, more underlying conditions may be detected should the existing condition be cleared.

When a test's operational state is anything other than up, the detection process will stop and any Detectable Tx Error fields will display a value of "none". History is not maintained for detectable transmit errors.

Not all errors are detectable. This function is only meant to guide the operator to a potential area of concern. There is a very large dependency on the direction of the MEP, the service type, the test protocol, and resource requirements that the test maintains over the underlying entity; any of these can influence the reporting of certain errors. Tests requiring the resources of an Up MEP, for example, will typically only report the conditions relating to incomplete configuration or the administrative state of the MEP. Up MEPs ignore the presence of egress connections and service states. However, LMM tests will report an unexpected error condition when the Up MEP SDP binding is down because it explicitly requires resources from the non-operational SDP binding. The same cannot be said if an LMM test was executing from an Up MEP configured over a non-operational SAP.

The TIMETRA-OAM-PM-MIB TEXTUAL-CONVENTION

TmnxOamPmDetectableTxError MIB lists all possible detectable error conditions.

The **show oam-pm session** *session-name* command provides a detailed view of the session, including each test and the Detectable Tx Err field for those tests.

The **show oam-pm sessions detectable-tx-errors** command lists all sessions that include a test with a detectable transmit error and the associated error.

3.10.4 Measurement Intervals

A measurement interval is a window of time that compartmentalizes the gathered measurements for an individual test that has 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 one per criteria per test.

Measurement intervals can be 5 minutes (**5-mins**), 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 (wall clock), with or without an optional offset. The boundary-type can be test-aligned, which means the start of the measurement interval coincides with the **no shutdown** of the test, for proactive tests. By default the start boundary is clocked 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, that 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 and not test aligned measurement intervals. Only offsets less than the measurement interval duration are allowed. [Table 9](#) provides some examples of the start times of each measurement interval.

Table 9 Measurement Intervals Start Time

Offset	5-min	15-min	1-hour	1-day
0 (default)	00,5,10,15..55	00,15,30,45	00 (top of the hour)	midnight
10 minutes	rejected	10,25,40,55	10 min after the hour	10 minutes after midnight
30 minutes	rejected	rejected	30 minutes after the hour	30 minutes after midnight
60 minutes	rejected	rejected	rejected	01:00am

Although test aligned approaches may seem beneficial for simplicity, there are some drawbacks that need to be considered. The goal of 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 will be maintained in volatile system memory by default. The number of intervals-stored is configurable per measurement interval. Different measurement interval lengths 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** can be used. If there is an active test, it can be viewed using the interval-number 1. In this case, the first completed record would be 2, previously completed would number back to the maximum intervals stored value plus one.

As new tests for the measurement interval complete, the older entries will get renumbered to maintain their relative position to the current test. As the retained test data for a measurement interval consumes the final entry, any subsequent entries will cause the removal of the oldest data.

There are obvious drawbacks to this storage model. Any high availability function that causes an active CPM switch will flush the results that were 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 write the completed data in system memory to non-volatile flash 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 the only necessary data be stored in non-volatile memory to avoid unacceptable risk and unnecessary resource consumption. It is also suggested that a large overlap between the data written to flash and stored in volatile memory is unnecessary.

The statistical information in system memory is also available by SNMP. If this method is chosen then a balance must be struck between the intervals retained and the times at which the SNMP queries collect the data. One must be cautious when determining the collection times through SNMP. If a file completes 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 10 OAM-PM XML Keywords and MIB Reference

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
oampm	—	None - header only

Table 10 OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
Keywords Shared by all OAM-PM Protocols		
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 in seconds	tmnxOamPmStsBaseElapsedTime
ftx	Frames sent	tmnxOamPmStsBaseTestFramesTx
frx	Frames received	tmnxOamPmStsBaseTestFramesRx
sus	Suspect flag	tmnxOamPmStsBaseSuspect
dmm	Delay Record	None - header only
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

Table 10 OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
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
mrbb	minimum frame delay range, backward	tmnxOamPmStsDelayDmmBwdMin
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
frr	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

Table 10 OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
lbo	Configured lower bound of the bin	tmnxOamPmCfgBinLowerBound
cnt	Number of measurements within the configured delay range. Note that the session_name, interval_duration, interval_number, {fd, fdr, ifdv}, bin_number, and {forward, backward, round-trip} indices are all 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
rxr	Received frames in the forward direction	tmnxOamPmStsLossSlmRxFwd
txb	Transmitted frames in the backward direction	tmnxOamPmStsLossSlmTxBwd
rxr	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
uaf	Undetermined available count in the forward direction	tmnxOamPmStsLossSlmUndtAvlFwd
uab	Undetermined available count in the backward direction	tmnxOamPmStsLossSlmUndtAvlBwd
uuf	Undetermined unavailable count in the forward direction	tmnxOamPmStsLossSlmUndtUnavlFwd
uub	Undetermined unavailable count in the backward direction	tmnxOamPmStsLossSlmUndtUnavlBwd
hlf	Count of HLIs in the forward direction	tmnxOamPmStsLossSlmHliFwd
hlf	Count of HLIs in the backward direction	tmnxOamPmStsLossSlmHliBwd
chf	Count of CHLIs in the forward direction	tmnxOamPmStsLossSlmChliFwd
chb	Count of CHLIs in the backward direction	tmnxOamPmStsLossSlmChliBwd
mff	minimum FLR in the forward direction	tmnxOamPmStsLossSlmMinFlrFwd
xff	maximum FLR in the forward direction	tmnxOamPmStsLossSlmMaxFlrFwd

Table 10 OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
aff	average FLR in the forward direction	tmnxOamPmStsLossSlmAvgFlrFwd
mfb	minimum FLR in the backward direction	tmnxOamPmStsLossSlmMinFlrBwd
xfb	maximum FLR in the backward direction	tmnxOamPmStsLossSlmMaxFlrBwd
afb	average FLR in the backward direction	tmnxOamPmStsLossSlmAvgFlrBwd
Imm	Frame Loss Measurement Record	None - header only
txf	Transmitted frames in the forward direction	tmnxOamPmStsLossLmmTxFwd
rxr	Received frames in the forward direction	tmnxOamPmStsLossLmmRxFwd
txb	Transmitted frames in the backward direction	tmnxOamPmStsLossLmmTxBwd
rxr	Received frames in the backward direction	tmnxOamPmStsLossLmmRxBwd
mff	minimum FLR in the forward direction	tmnxOamPmStsLossLmmMinFlrFwd
xff	maximum FLR in the forward direction	tmnxOamPmStsLossLmmMaxFlrFwd
aff	average FLR in the forward direction	tmnxOamPmStsLossLmmAvgFlrFwd
mfb	minimum FLR in the backward direction	tmnxOamPmStsLossLmmMinFlrBwd
xfb	maximum FLR in the backward direction	tmnxOamPmStsLossLmmMaxFlrBwd
afb	average FLR in the backward direction	tmnxOamPmStsLossLmmAvgFlrBwd
ave	Imm availability enabled/disabled	No TIMETRA-OAM-PM-MIB entry
avf	available count in the forward direction	tmnxOamPmStsLossLmmAvailIndFwd
avb	available count in the backward direction	tmnxOamPmStsLossLmmAvailIndBwd
uvf	unavailable count in the forward direction	tmnxOamPmStsLossLmmUnavIndFwd
uvb	unavailable count in the backward direction	tmnxOamPmStsLossLmmUnavIndBwd
uaf	undetermined available count in the forward direction	tmnxOamPmStsLossLmmUndtAviFwd
uab	undetermined available count in the backward direction	tmnxOamPmStsLossLmmUndtAviBwd
uuf	undetermined unavailable count in the forward direction	tmnxOamPmStsLossLmmUndtUnavIFwd
uub	undetermined unavailable count in the backward direction	tmnxOamPmStsLossLmmUndtUnavIBwd

Table 10 OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
hlf	count of HLIs in the forward direction	tmnxOamPmStsLossLmmHliFwd
hlb	count of HLIs in the backward direction	tmnxOamPmStsLossLmmHliBwd
chf	count of CHLIs in the forward direction	tmnxOamPmStsLossLmmChliFwd
chb	count of CHLIs in the backward direction	tmnxOamPmStsLossLmmChliBwd
udf	undetermined delta-t in the forward direction	tmnxOamPmStsLossLmmUndetDelTsFwd
udb	undetermined delta-t in the backward direction	tmnxOamPmStsLossLmmUndetDelTsBwd
TLD	TWAMP Light Delay Record	None - header only
mdr	minimum frame delay, round-trip	tmnxOamPmStsDelayTwl2wyMin
xdr	maximum frame delay, round-trip	tmnxOamPmStsDelayTwl2wyMax
adr	average frame delay, round-trip	tmnxOamPmStsDelayTwl2wyAvg
mdf	minimum frame delay, forward	tmnxOamPmStsDelayTwlFwdMin
xdf	maximum frame delay, forward	tmnxOamPmStsDelayTwlFwdMax
adf	average frame delay, forward	tmnxOamPmStsDelayTwlFwdAvg
mdb	minimum frame delay, backward	tmnxOamPmStsDelayTwlBwdMin
xdb	maximum frame delay, backward	tmnxOamPmStsDelayTwlBwdMax
adb	average frame delay, backward	tmnxOamPmStsDelayTwlBwdAvg
mvr	minimum inter-frame delay variation, round-trip	tmnxOamPmStsDelayTwl2wyMin
xvr	maximum inter-frame delay variation, round-trip	tmnxOamPmStsDelayTwl2wyMax
avr	average inter-frame delay variation, round-trip	tmnxOamPmStsDelayTwl2wyAvg
mvf	minimum inter-frame delay variation, forward	tmnxOamPmStsDelayTwlFwdMin
xvf	maximum inter-frame delay variation, forward	tmnxOamPmStsDelayTwlFwdMax
avf	average inter-frame delay variation, forward	tmnxOamPmStsDelayTwlFwdAvg
mvb	minimum inter-frame delay variation, backward	tmnxOamPmStsDelayTwlBwdMin

Table 10 OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
xvb	maximum inter-frame delay variation, backward	tmnxOamPmStsDelayTwIBwdMax
avb	average inter-frame delay variation, backward	tmnxOamPmStsDelayTwIBwdAvg
mrr	minimum frame delay range, round-trip	tmnxOamPmStsDelayTwl2wyMin
xrr	maximum frame delay range, round-trip	tmnxOamPmStsDelayTwl2wyMax
arr	average frame delay range, round-trip	tmnxOamPmStsDelayTwl2wyAvg
mrf	minimum frame delay range, forward	tmnxOamPmStsDelayTwlFwdMin
xrf	maximum frame delay range, forward	tmnxOamPmStsDelayTwlFwdMax
arf	average frame delay range, forward	tmnxOamPmStsDelayTwlFwdAvg
mrbb	minimum frame delay range, backward	tmnxOamPmStsDelayTwlBwdMin
xrbb	maximum frame delay range, backward	tmnxOamPmStsDelayTwlBwdMax
arbb	average frame delay range, backward	tmnxOamPmStsDelayTwlBwdAvg
fdr	frame delay bin record, round-trip	None - header only
fdf	frame delay bin record, forward	None - header only
fddb	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
frr	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

Table 10 OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
cnt	Number of measurements within the configured delay range. Note that the session_name, interval_duration, interval_number, {fd, fdr, ifdv}, bin_number, and {forward, backward, round-trip} indices are all provided by the surrounding XML context.	tmnxOamPmStsDelayTwlBinFwdCount tmnxOamPmStsDelayTwlBinBwdCount tmnxOamPmStsDelayTwlBin2wyCount
TLL	TWAMP Light Loss Record	None - header only
slm	Synthetic Loss Measurement Record	None - header only
txf	Transmitted frames in the forward direction	tmnxOamPmStsLossTwlTxFwd
rxfr	Received frames in the forward direction	tmnxOamPmStsLossTwlRxFwd
txb	Transmitted frames in the backward direction	tmnxOamPmStsLossTwlTxBwd
rxbr	Received frames in the backward direction	tmnxOamPmStsLossTwlRxBwd
avf	Available count in the forward direction	tmnxOamPmStsLossTwlAvailIndFwd
avbr	Available count in the backward direction	tmnxOamPmStsLossTwlAvailIndBwd
uvf	Unavailable count in the forward direction	tmnxOamPmStsLossTwlUnavlIndFwd
uvbr	Unavailable count in the backward direction	tmnxOamPmStsLossTwlUnavlIndBwd
uaf	Undetermined available count in the forward direction	tmnxOamPmStsLossTwlUndtAvlFwd
uabr	Undetermined available count in the backward direction	tmnxOamPmStsLossTwlUndtAvlBwd
uuf	Undetermined unavailable count in the forward direction	tmnxOamPmStsLossTwlUndtUnavlFwd
uubr	Undetermined unavailable count in the backward direction	tmnxOamPmStsLossTwlUndtUnavlBwd
hlf	Count of HLIs in the forward direction	tmnxOamPmStsLossTwlHliFwd
hlbr	Count of HLIs in the backward direction	tmnxOamPmStsLossTwlHliBwd
chf	Count of CHLIs in the forward direction	tmnxOamPmStsLossTwlChliFwd
chlbr	Count of CHLIs in the backward direction	tmnxOamPmStsLossTwlChliBwd
mff	minimum FLR in the forward direction	tmnxOamPmStsLossTwlMinFlrFwd
xff	maximum FLR in the forward direction	tmnxOamPmStsLossTwlMaxFlrFwd

Table 10 OAM-PM XML Keywords and MIB Reference (Continued)

XML File Keyword	Description	TIMETRA-OAM-PM-MIB Object
aff	average FLR in the forward direction	tmnxOamPmStsLossTwiAvgFlrFwd
mfb	minimum FLR in the backward direction	tmnxOamPmStsLossTwiMinFlrBwd
xfb	maximum FLR in the backward direction	tmnxOamPmStsLossTwiMaxFlrBwd
afb	average FLR in the backward direction	tmnxOamPmStsLossTwiAvgFlrBwd

By default, 5-mins measurement interval will store 33 test runs (32+1) with a configurable range of [1 to 96]. By default, 15-mins measurement interval will store 33 test runs (32+1) with a configurable range of [1 to 96]. The 5-mins and 15-mins measurement intervals share the [1 to 96] pool up to a maximum of 96. In the unlikely case where both the 5-mins and 15-mins measurement intervals are configured for the same oam-pm session, the total combined intervals stored cannot exceed 96. By default, 1-hour measurement intervals will store 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). When the 1-day measurement interval is configured, this is the only value for intervals. The value cannot be changed.

All four measurement intervals may included for a single session if required. Each measurement interval that is included in a session will be updated simultaneously for each test that is being executed. If a measurement interval duration is not required, it should not be configured. In addition to the four predefined lengths, a fifth measurement interval is always on and is allocated at test creation, the “raw” measurement interval. 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 in order to flush stale statistical data in order to look at current conditions. This measurement interval has no configuration options, and it cannot be written to flash and cannot be disabled. It is a single never ending collection window.

Memory allocation for the measurement intervals is performed when the test transitions from an operationally down state to an operationally up state. Any previous stored test data will be cleared from volatile memory in favor of the new allocation. This will result in the loss of all data that has not been written to the XML file or collected by some other means. Volatile memory will be flushed and completely released when the test is deleted from the configuration, or a high availability event causes the backup CPM to become the newly active CPM, or some other event clears the active CPM system memory. Following an HA event, memory reallocation occurs when the operational state of the test changes from down to up. 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 to true when there are times of service disruption, maintenance windows, discontinuity, low packet counts, or other such type events. Higher level systems would be required to interpret and correlate those types of event for measurement intervals that are 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 to those higher level system to discount the specific windows of time.

3.10.5 Data Structures and Storage

There are two main metrics that are the focus of OAM-PM, delay and loss. The different metrics have their own unique storage structures and will allocate their own measurement intervals for these structures. This is regardless of whether the performance data is gathered with a single packet or multiple packet types.

Delay metrics include the following:

- Frame Delay (FD)- The amount of time it takes to travel form the source to the destination and back
- InterFrame Delay Variation (IFDV) -The difference in the delay metrics between two adjacent packets
- Frame Delay Range (FDR)-The difference between the minimum frame delay and the individual packet
- Mean Frame Delay (MFD) -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 a min/max/average

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 one 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, consideration must be given to Precision Time Protocol (PTP) 1588v2.

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 bin 0, the lowest bin which has a lower boundary of 0 microseconds. In order to isolate these outlying negative results, the lower boundary of bin 1 for the frame delay type could be set to a value of 1 micro second. This means bin 0 would then only collect results that are 1 micro second or less. This would be an indication of the number of negative results that are being collected.

Delay results are mapped to the measurement interval that is active when the result arrives back at the source.

Loss metrics are only unidirectional and will report Frame Loss Ratio (FLR) and availability information. Frame loss ratio is the percentage computation of loss (lost/sent). 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 loss or received based on sequence numbers in the protocol. A number of probes are rolled up into a small measurement window (delta-t). 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 high loss. If the resulting percentage is lower than or equal to the threshold, the small window is marked as non-high loss. A sliding window is defined as some number of small windows. 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. The maximum size of the sliding window cannot be greater than 100 seconds. The default values for these availability parameters can differ from PDU type to PDU type.

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 reporting 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 periods of complete or extended failure, 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 will maintain the previous state updating the appropriate directional availability or unavailability counter. At the same time, an additional per direction undetermined counter will be updated. This undetermined counter is used to indicate that the availability or unavailability statistics were indeterminable for a number of small windows.

During connectivity outages the higher level systems could 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 closing a measurement interval could be equal to the length of the sliding window and the timeout of the last packet. A measurement interval cannot be closed 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, standards bodies indicate the timeout value per packet. For Ethernet, the timeout value for DMMv1, LMM, and SLM is set at 5s and is not configurable.

There are no log events based on availability or unavailability state changes. Based on the subjective nature of these counters, considering complete failure or total loss when it may not be possible to determine availability or unavailability, these counters represent the raw values that must be interpreted.

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 small windows is one second in length. A CHLI is consecutive high loss intervals that exceeds a consecutive threshold within the sliding window. Only one CHLI will be counted within a window. HLI and CHLI counters are only incremented during periods of availability. These counters are not incremented during periods of unavailability.

Both ETH-SLM and ETH-LMM provide methods for reporting loss. ETH-SLM uses the synthetic packets on the wire. ETH-LMM monitors the amount of service data. ETH-LMM frame loss counting is significantly different from ETH-SLM synthetic packet counting. The PDUs provide the largest variance.

The SLM PDU includes a sequence number that allows the mapping of message to the appropriate responses. The LMM PDU is fixed without support for optional TLVs. ETH-LMM has no method of correlating message and response. This means the LMR could represent a sample window equating to more than one LMM. SLM produces known and constant load on the network, whereas service data frames, counted by LMM, will vary or possibly be null. These two facts require slightly different approaches to availability and reliability.

LMM requires the **availability** command to not be shut down in order to enable the collection and computation of availability. When availability is not enabled, the existing behavior for frame loss ratio min/max/avg remains unchanged. Determining if any new min/max has been encountered and computing the avg is based on every individual LMR. When LMM availability is enabled, the determination of the new min/max is based on the delta-t window. The longer sample size reduces the impact of gain and loss that may be introduced by reordering. No new FLR min/max values will be considered during periods of determined unavailability, regardless of the configuration for the availability option. The average FLR computation depends on the configuration of the availability metric. If availability is enabled, FLR is based on the sum of all LMR calculations received during periods of availability. If availability is not enabled, FLR is based on the LMR received prior to the closing of the measurement interval.


LMM will continue to count and report transmit and receive delta frames collected by the **collect-lmm-stats** command for the SAP, regardless of availability state. LMM will continue to display data frames counted during periods of unavailability; however, these frames will not count towards the average FLR of the measurement interval. In contrast, SLM does not count the synthetic packets on the wire during periods of unavailability. The raw transmit and receive information gathered by LMM has significant value regardless of availability in an unavailability state.

LMM includes a new counter, undetermined-delta-t, for the forward and backward directions. This new counter counts the number of delta-t windows that have no LMR responses recorded in that window of time, and provide an indication of the quality and scope of the delta-ts.

[Figure 42](#) looks at loss in a single direction using synthetic packets. It demonstrates what happens when a possible unavailability event crosses a measurement interval boundary. As shown, the first 13 small windows are all marked available (1). This means that the lost 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. This means that the lost 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 (FLR) 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 an unavailability event, they would be marked as HLI or CHLI and be counted as part of the frame loss ratio.

Figure 42 Evaluating and Computing Loss and Availability



al_0387

3.10.6 Bin Groups

Bin groups are templates that are referenced by the session. Three types of binnable statistics are available:

- Frame Delay (FD); round trip, forward and backward
- InterFrame Delay Variation (IFDV); round trip, forward and backward
- Frame Delay Range (FDR); round trip, forward and backward

Each of these metrics can have up to 10 bins configured to group the results. Bins are configured by indicating a lower boundary. Bin 0 has a lower boundary that is always zero and not configurable. The micro second range of the bins is the difference between the adjacent lower boundaries. For example, bin-type fd bin 1 configured with a lower-bound 1000 micro seconds means bin 0 will capture all frame delay statistics results between 0 and 1ms. Bin 1 will capture all results above 1ms 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.

A bin group configuration may relegate the first and last bin to capture anomalous measurements. Anomalous measurements can result from legacy equipment that queues a number of packets prior to circuit establishment. The relegation model characterizes these bins as non-representative of real network delay measurements. Results in these bins should be omitted from the average (avg) calculation. To accommodate these models, the command **exclude-from-avg** is available under the **config>oam-pm>bin-group>bin-type** hierarchy. This excludes results from the rolling average calculation that map to the excluded bins. The bins statistics will still accumulate all of the results, but the results are not part of the average computation. Every configured bin is included in the average calculation by default.

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 or the values that meet the desired goals.

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 shutdown. If there is a requirement to change the setting of a bin group where a large number of sessions are referencing a bin group, migrating existing sessions to a new bin group with the new parameters could be considered to reduce the maintenance window.

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. Any bin lower-bound value that aligns with the 5000 μ s (5 ms) default value (bin number $\times$ 5000 μ s) will not be displayed as part of the output of the **info** command within the configuration. The **info** command does not display default values. The **info detail** command is required to show the default values. The bin group 1 configuration parameters are shown below.

```
-----
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	-	-

3.10.7 Relating the Components

Figure 43 brings together all the concepts discussed in the OAM-PM architecture. It shows a more detailed hierarchy than previously shown in the introduction including the relationship between the tests, the measurement intervals, and the storage of the results.

Figure 43 is a logical representation and not meant to represent the exact flow between elements in the architecture. For example, the line connecting the "Acct-Policy" and the "Intervals Stored & Collected" is not intended to show the accounting policy being responsible for the movement of data from completed records "to Be Collected" to "Collected".

Figure 43 Relating OAM-PM Components



3.10.8 IP Performance Monitoring

The following configuration includes the configuration information specific to the TWAMP Light session controller. It does not include the MPLS configuration or the TWAMP Light session-responder configuration. For complete details on configuring the session-responder, refer to “TWAMP Light” in the IP Performance Monitoring (IP PM) section.

3.10.8.1 Accounting Policy Configuration

```
config>log# info
-----
      file-id 2
        description "IP OAM PM XML file Paramaters"
        location cf2:
        rollover 15 retention 2
      exit
    accounting-policy 2
      description "IP OAM PM Collection Policy for 15-min MI"
      record complete-pm
      collection-interval 10
      to file 2
      no shutdown
    exit
  log-id 1
  exit
-----
```

3.10.8.2 Service Configuration

```
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 64364 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
```

3.10.8.3 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
        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 "ip-vprn-500" test-family ip session-type proactive create
    bin-group 2
    no description
    meas-interval 15-mins create
        accounting-policy 2
        boundary-type clock-aligned
        clock-offset 0
        intervals-stored 8
    exit
    ip
        dest-udp-port 64364
        destination 10.1.1.1
        fc "l2"
        no forwarding
        profile in
        router 500
        source 10.2.1.1
        ttl 255
        twamp-light test-id 500 create
            interval 1000
            loss
                flr-threshold 50
                timing frames-per-delta-t 10 consec-delta-t 10 chli-
threshold 5
                exit
            pad-size 27
            record-stats delay-and-loss
            no test-duration
            no shutdown
        exit
    exit

```

3.10.9 Ethernet Performance Monitoring

The following configuration provides an example comprised of the different Ethernet OAM PM elements using ETH-CFM tools.

3.10.9.1 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
-----
```

3.10.9.2 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
-----
```

3.10.9.3 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
sap 1/2/1:4.* create
exit
no shutdown
```

3.10.9.4 Ethernet 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
  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

```

The RAW measurement interval can also use the monitor command to automatically update the statistics.

The following configuration and show commands provide an example of how frame loss measurement (ETH-LMM) can be used to collect frame loss metrics and the statistics gathered.

The LMM reflector must be configured to collect the statistics on the SAP or MPLS SDP binding where the terminating MEP has been configured.

```
epipe 1000 customer 1 create
    sap 1/1/10:1000.* create
    exit
    spoke-sdp 1:1000 create
        eth-cfm
        collect-lmm-stats
        mep 31 domain 14 association 1000 direction down
        ccm-enable
        mac-address 00:00:00:00:00:31
        no shutdown
    exit
    exit
    no shutdown
    exit
    no shutdown
    exit
```

The launch point must also enable statistical collection on the SAP or MPLS SDP binding of the MEP launch point.

```
epipe 1000 customer 1 create
    sap 1/1/10:1000.* create
    exit
    spoke-sdp 1:1000 create
        eth-cfm
        collect-lmm-stats
        mep 28 domain 14 association 1000 direction down
        no shutdown
    exit
    exit
    no shutdown
    exit
    no shutdown
    exit
```

The launch point must configure the OAM-PM session parameters. The CLI below shows a session configured with DMM for delay measurements (1s intervals) and LMM for frame loss measurements (10s interval). When using LMM for frame loss, the frame loss ratio and the raw frame transmit and receive statistics are captured, along with basic measurement interval and protocol information.

```

session "eth-pm-service-1000" test-family ethernet session-type proactive create
    bin-group 2
    description "Frame Loss using LMM"
    meas-interval 15-mins create
        accounting-policy 2
        intervals-stored 8
    exit
    ethernet
        dest-mac 00:00:00:00:00:31
        source mep 28 domain 14 association 1000
        dmm test-id 1000 create
            no shutdown
        exit
        lmm test-id 1000 create
            interval 10000
            no shutdown
        exit
    exit
exit
-----

```

3.10.10 OAM-PM Event Monitoring

The previous section described the OAM-PM architecture. That provides a very powerful and well defined mechanism to collect key performance information. This data is typically uploaded to higher level systems for consolidation and reporting tracking performance trends and conformance to Service Level Agreements (SLA). Event monitoring (**event-mon**) allows thresholds to be applied to the well defined counters, percentage and binned results for a single and measurement interval per session. This Traffic Crossing Alert (TCA) function can be used to raise a log event when a configured threshold is reached. Optionally, The TCA can be cleared if a clear threshold is not breached in a subsequent measurement interval.

Thresholds can be applied to binned delay metrics and the various loss metric counters or percentages. The type of the TCA is based on the configuration of the two threshold values, **threshold** *raise-threshold* and **clear** *clear-threshold*. The on network element TCA functions are provided to log an event that is considered an exception condition that requires immediate attention. A single threshold can be applied to the collected metric.

Stateless TCAs are those events that do not include a configured *clear-threshold*. Stateless TCAs will raise the event when the *raise-threshold* is reached but do not share state with any following measurement intervals. Each subsequent measurement interval is treated as a unique entity without previous knowledge of any alerts raised. Each measurement interval will consider only its data collection and raise all TCAs as the thresholds are reached. A stateless event raised in one measurement interval silently expires at the end of that measurement interval without an explicit clear event.

Stateful TCAs require the configuration of the optional **clear** *clear-threshold*. Stateful TCAs will raise the event when the *raise-threshold* is reached and carry that state forward to subsequent measurement intervals. That state is maintained and no further raise events will be generated for that monitored event until a subsequent measurement interval completes and the value specified by the clear-threshold is not reached. When a subsequent measurement interval completes and the specific *clear-threshold* is not crossed an explicit clear log event is generated. Clear events support a value of zero which means that the event being cleared must have no errors at the completion of the measurement interval to clear a previous raise event. At this point, the event is considered cleared and a raise is possible when the next **threshold** *raise-threshold* is reached.

The raise threshold must be higher than the clear threshold. The only time both can be equal is if they are disabled.

Alerts can only be raised and cleared once per measurement interval per threshold. Once a raise is issued no further monitoring for that event occurs in that measurement interval. A clear is only logged at the end of a subsequent measurement interval following a raise and only for stateful event monitoring.

Changing threshold values or events to monitor for the measurement interval do not require the individual tests within the session or the related resource (**bin-group**) to be shutdown. Starting the monitoring process, adding a new event to monitor, or altering a threshold will stop the existing function that has changed with the new parameters activated at the start of the next measurement interval. Stopping the monitoring or removing an event will maintain the current state until the completion of the adjacent measurement interval after which any existing state will be cleared.

OAM-PM sessions may have multiple measurement intervals. Event monitoring can only be configured against a single configured measurement interval per session.

Delay event thresholds can be applied to Frame Delay (FD), InterFrame Delay Variation (IFDV) and Frame Delay Range (FDR). These are binned delay metrics with directionality, forward, backward and round-trip. Configuration of event thresholds for these metrics are within the **config>oam-pm>bin-group** *bin-group-number* and applied to a specific bin-type. The **delay-event** specifies the direction that is to be measured {**forward** | **backward** | **round-trip**}, the thresholds and the lowest bin number. The lowest bin value applies the threshold to the cumulative results in that bin and all higher. The default bin group (bin-group 1) cannot be modified and as such does not support the configuration of event thresholds. A session that makes use of a bin group inherits those bin group attributes including delay event threshold settings.

If the operator subscribes to a model that relegates one or more of the highest bins to anomalous results, these results should not be included in the TCA count. The **delay-event-exclusion** command is available under the **config>oam-pm>bin-group>bin-type** hierarchy. This command will exclude any results in the specified bin, along with the results in any bin higher than the one specified, from the TCA count. In order to use this command, a **delay-event** in the same direction for the same bin type must be configured. The excluded bins must be higher than the TCA threshold configured using the **threshold raise-threshold** command. This command is similar to the **delay-event** command. This does not require the bin group to be shut down. On-the-fly changes will cause the delay event to be suspended until the next measurement interval for the affected bin type.

Ethernet supports gathering delay information using the ETH-DMM protocol. IP supports the gathering of delay information using the TWAMP Light function.

Loss events and threshold are configured within the session under the specific loss based protocol. Loss event thresholds can be applied to the average Frame Loss Ratio (FLR) in the forward and backward direction. This event is analyzed at the end of the measurement interval to see if the computed FLR is equal to or higher than the configured threshold as a percentage. The availability and reliability loss events may be configured against the counts in forward and backward direction as well as the aggregate (sum of both directions). The aggregate is only computed for thresholds and not stored as an independent value in the standard OAM-PM loss dataset. The availability and reliability loss events include the high loss interval (HLI), Consecutive HLI (CHLI), unavailability, undetermined availability and undetermined unavailability.

Ethernet supports the gathering of loss information using ETH-SLM and ETH-LMM. IP supports the gathering of loss information using TWAMP Light functionality. ETH-SLM, ETH-LMM, and TWAMP Light support threshold configuration for FLR and the availability and reliability loss events.

Configuring the event threshold and their behavior, stateless or stateful, completes the first part of the requirement. The event monitoring function must be enabled per major function, delay or loss. This is configured under the measurement interval that is used to track events. One measurement interval per session can be configured to track events. If event tracking of type, delay or loss, is configured against a measurement interval within the session no other measurement interval can be used to track events. For example, if the measurement interval 15-min for oam-pm session eth-pm-session1 has delay-events active, no other measurement interval within that session can be used to track delay or loss-events.

When a raise threshold is reached a log event warning is generated from the OAM application using the number 2300. If the event is stateful, **clear clear-threshold** configured, an explicit clear will be logged when a subsequent measurement interval does not exceed the clear threshold. The clear event is also a warning message from the OAM protocol but uses number 2301.

The session name is included as part of the subject.

A more detailed message is included immediately following the subject. This includes:

- type of the event — raised or cleared
- session name in quotations
- test type — representing the protocol (dmm, slm, lmm, or twl)
- the start time of the measurement interval in UTC format
- delay bin type — fd, fdr, ifdv. Not applicable for loss measurement.
- threshold type — the metric type that is covered by this alarm. Delay will include the various delay metrics, and loss will include the various loss metrics.
- direction — forward, backward, round-trip or aggregate
- bin lower bound (us) — the lower bound value of the lowest bin associated with the TCA. Not applicable for loss measurement.
- configured threshold — the value of the threshold
- operational value — the measured value relating to the action
- tca type — stateful or stateless
- suspect flag — copied from the measurement interval (events do not affect the suspect flag)

```
25 2016/01/06 15:45:15.42 UTC WARNING: OAM #2301 Base vpls1000-PM-YL4-1/1/9:1000.1000
```

```
"OAM-PM TCA cleared for session "vpls1000-PM-YL4-1/1/9:1000.1000", test type slm, measurement interval duration 5-mins, MI start 2016/01/
```

```
06 15:40:00 UTC, delay bin type
```

```
not-applicable. Threshold type loss-hli, direction aggregate, bin lower bound (us) not-applicable, configured threshold 0, operational value 0. TCA type stateful, suspect flag false."
```

```
24 2016/01/06 15:39:04.42 UTC WARNING: OAM #2300 Base vpls1000-PM-YL4-1/1/9:1000.1000
```

```
"OAM-PM TCA raised for session "vpls1000-PM-YL4-1/1/9:1000.1000", test type slm, measurement interval duration 5-mins, MI start 2016/01/
```

```
06 15:35:00 UTC, delay bin type
```

```
not-applicable. Threshold type loss-hli, direction aggregate, bin lower bound (us) not-applicable, configured threshold 50, operational value 50. TCA type stateful, suspect flag false."
```

```
3 2016/01/06 09:10:00.00 UTC WARNING: OAM #2301 Base vpls1000-PM-YL4-1/1/9:1000.1000
```

```
"OAM-PM TCA cleared for session "vpls1000-PM-YL4-1/1/9:1000.1000", test type dmm, measurement interval duration 5-mins, MI start 2016/01/
```

```
06 09:05:00 UTC, delay bin type
```

```
fd. Threshold type delay, direction round-trip, bin lower bound (us) 2000, configured threshold 10, operational value 10. TCA type stateful, suspect flag false."
```

```
2 2016/01/06 08:09:47.68 UTC WARNING: OAM #2300 Base vpls1000-PM-YL4-1/1/9:1000.1000
```

```
"OAM-PM TCA raised for session "vpls1000-PM-YL4-1/1/9:1000.1000", test type dmm, measurement interval duration 5-mins, MI start 2016/01/
```

```
06 08:05:00 UTC, delay bin type
fd. Threshold type delay, direction round-trip, bin lower bound (us) 2000,
configured
threshold 50, operational value 50. TCA type stateful, suspect flag false."
```

Only those events deemed important should be configured and activated per session.

A simple Ethernet session example is provided to show the basic configuration and monitoring of threshold event monitoring.

The bin group is configured for the required thresholds.

```
bin-group 4 fd-bin-count 10 fdr-bin-count 2 ifdv-bin-count 10 create
  bin-type fd
    bin 1
      lower-bound 1
    exit
    bin 2
      lower-bound 1000
    exit
    bin 3
      lower-bound 2000
    exit
    bin 4
      lower-bound 3000
    exit
    bin 5
      lower-bound 4000
    exit
    bin 6
      lower-bound 5000
    exit
    bin 7
      lower-bound 6000
    exit
    bin 8
      lower-bound 7000
    exit
    bin 9
      lower-bound 8000
    exit
    delay-event round-trip lowest-bin 6 threshold 10
  exit
  bin-type ifdv
    bin 1
      lower-bound 200
    exit
    bin 2
      lower-bound 400
    exit
    bin 3
      lower-bound 600
    exit
    bin 4
      lower-bound 800
    exit
```

```

        bin 5
            lower-bound 1000
        exit
        bin 6
            lower-bound 1200
        exit
        bin 7
            lower-bound 1400
        exit
        bin 8
            lower-bound 1600
        exit
        bin 9
            lower-bound 1800
        exit
        delay-event round-trip lowest-bin 7 threshold 30 clear 20
    exit
    no shutdown
exit

```

The OAM-PM session contains all the session attributes, test attributes and the loss event thresholds and the configuration of the event monitoring functions.

```

session "eth-pm-service-1100" test-family ethernet session-type proactive create
    bin-group 4
    description "Service 1000 PM Collection"
    meas-interval 15-mins create
        accounting-policy 2
        event-mon
            delay-events
            loss-events
            no shutdown
        exit
        intervals-stored 8
    exit
    ethernet
        dest-mac 00:00:00:00:00:31
        source mep 28 domain 15 association 1000
        dmm test-id 1000 create
            no shutdown
        exit
        slm test-id 1000 create
            loss-events
                avg-flr-event forward threshold 2.000
                avg-flr-event backward threshold 2.000
                hli-event aggregate threshold 27 clear 9
            exit
            timing frames-per-delta-t 1 consec-delta-t 10 interval 1000 chli-
threshold 5
            no shutdown
        exit
    exit
exit

```

3.11 Traceroute with ICMP Tunneling In Common Applications


This section provides sample output of the traceroute OAM tool when the ICMP tunneling feature is enabled in a few common applications.

The ICMP tunneling feature is described in [Tunneling of ICMP Reply Packets over MPLS LSP](#) and provides supports for appending to the ICMP reply of type Time Exceeded the MPLS label stack object defined in RFC 4950. The new MPLS Label Stack object permits an LSR to include label stack information including label value, EXP, and TTL field values, from the encapsulation header of the packet that expired at the LSR node.

This section describes the following topics:

- [BGP-LDP Stitching and ASBR/ABR/Data Path RR for BGP IPv4 Label Route](#)
- [VPRN Inter-AS Option B](#)
- [VPRN Inter-AS Option C and ASBR/ABR/Data Path RR for BGP IPv4 Label Route](#)

3.11.1 BGP-LDP Stitching and ASBR/ABR/Data Path RR for BGP IPv4 Label Route



```
# lsp-trace ldp-bgp stitching
*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 detail downstream-map-
tlv ddmap
lsp-trace to 10.20.1.6/32: 0 hops min, 0 hops max, 108 byte packets
1 10.20.1.1 rtt=2.89ms rc=15(LabelSwitchedWithFecChange) rsc=1
   DS 1: ipaddr=10.10.1.2 ifaddr=10.10.1.2 iftype=ipv4Numbered MRU=1496
        label[1]=262143 protocol=3(LDP)
        label[2]=262139 protocol=2(BGP)
        fecchange[1]=POP fecctype=LDP IPv4 prefix=10.20.1.6 remotepeer=0.0.0.0 (U
nknown)
        fecchange[2]=PUSH fecctype=BGP IPv4 prefix=10.20.1.6 remotepeer=10.20.1.2
        fecchange[3]=PUSH fecctype=LDP IPv4 prefix=10.20.1.2 remotepeer=10.10.1.2
```

```

2 10.20.1.2 rtt=5.19ms rc=3(EgressRtr) rsc=2
2 10.20.1.2 rtt=5.66ms rc=8(DSRtrMatchLabel) rsc=1
    DS 1: ipaddr=10.10.4.4 ifaddr=10.10.4.4 iftype=ipv4Numbered MRU=0
        label[1]=262138 protocol=2(BGP)
3 10.20.1.4 rtt=6.53ms rc=15(LabelSwitchedWithFecChange) rsc=1
    DS 1: ipaddr=10.10.6.5 ifaddr=10.10.6.5 iftype=ipv4Numbered MRU=1496
        label[1]=262143 protocol=3(LDP)
        label[2]=262138 protocol=2(BGP)
        fecchange[1]=PUSH fectype=LDP IPv4 prefix=10.20.1.5 remotepeer=10.10.6.5
4 10.20.1.5 rtt=8.51ms rc=3(EgressRtr) rsc=2
4 10.20.1.5 rtt=8.45ms rc=15(LabelSwitchedWithFecChange) rsc=1
    DS 1: ipaddr=10.10.10.6 ifaddr=10.10.10.6 iftype=ipv4Numbered MRU=1496
        label[1]=262143 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=11.2ms rc=3(EgressRtr) rsc=1

```

```

*A:Dut-A# configure router ldp-
shortcut (to add ldp label on first hop but overall behavior is similar)

```

12.0R4 default behavior (we have routes back to the source)

```

*A:Dut-A# traceroute 10.20.1.6 detail wait 100
traceroute to 10.20.1.6, 30 hops max, 40 byte packets
 1  1 10.10.2.1 (10.10.2.1) 3.47 ms
 1  2 10.10.2.1 (10.10.2.1) 3.65 ms
 1  3 10.10.2.1 (10.10.2.1) 3.46 ms
 2  1 10.10.1.2 (10.10.1.2) 5.46 ms
 2  2 10.10.1.2 (10.10.1.2) 5.83 ms
 2  3 10.10.1.2 (10.10.1.2) 5.20 ms
 3  1 10.10.4.4 (10.10.4.4) 8.55 ms
 3  2 10.10.4.4 (10.10.4.4) 7.45 ms
 3  3 10.10.4.4 (10.10.4.4) 7.29 ms
 4  1 10.10.6.5 (10.10.6.5) 9.67 ms
 4  2 10.10.6.5 (10.10.6.5) 10.1 ms
 4  3 10.10.6.5 (10.10.6.5) 10.9 ms
 5  1 10.20.1.6 (10.20.1.6) 11.5 ms
 5  2 10.20.1.6 (10.20.1.6) 11.1 ms
 5  3 10.20.1.6 (10.20.1.6) 11.4 ms

```

Enable ICMP tunneling on PE and ASBR nodes.

```

*A:Dut-D# # configure router ttl-propagate label-route-local all *A:Dut-
C,D,E,F# configure router icmp-tunneling

```

```

*A:Dut-C# traceroute 10.20.1.6 detail wait 100
traceroute to 10.20.1.6, 30 hops max, 40 byte packets
 1  1 10.10.1.1 (10.10.1.1) 11.8 ms
    returned MPLS Label Stack Object
    entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
 1  2 10.10.1.1 (10.10.1.1) 12.5 ms
    returned MPLS Label Stack Object
    entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
 1  3 10.10.1.1 (10.10.1.1) 12.9 ms
    returned MPLS Label Stack Object
    entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
 2  1 10.10.4.2 (10.10.4.2) 13.0 ms
    returned MPLS Label Stack Object
    entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0

```

```

        entry 2: MPLS Label = 262139, Exp = 7, TTL = 1, S = 1
2  2  10.10.4.2 (10.10.4.2) 13.0 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
        entry 2: MPLS Label = 262139, Exp = 7, TTL = 1, S = 1
2  3  10.10.4.2 (10.10.4.2) 12.8 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
        entry 2: MPLS Label = 262139, Exp = 7, TTL = 1, S = 1
3  1  10.10.6.4 (10.10.6.4) 10.1 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
3  2  10.10.6.4 (10.10.6.4) 11.1 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
3  3  10.10.6.4 (10.10.6.4) 9.70 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
4  1  10.10.10.5 (10.10.10.5) 12.5 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 7, TTL = 255, S = 0
        entry 2: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
4  2  10.10.10.5 (10.10.10.5) 11.9 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 7, TTL = 255, S = 0
        entry 2: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
4  3  10.10.10.5 (10.10.10.5) 11.8 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 7, TTL = 255, S = 0
        entry 2: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
5  1  10.20.1.6 (10.20.1.6) 12.2 ms
5  2  10.20.1.6 (10.20.1.6) 12.5 ms
5  3  10.20.1.6 (10.20.1.6) 13.2 ms

```

```

# With lsr-label-route all on all LSRs (only needed on Dut-E) *A:Dut-
E# configure router ttl-propagate lsr-label-route all

```

```

*A:Dut-

```

```

A# traceroute 10.20.1.6 detail wait 100 traceroute to 10.20.1.6, 30 hops max, 40 byte packets

```

```

1  1  10.10.1.1 (10.10.1.1) 12.4 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
1  2  10.10.1.1 (10.10.1.1) 11.9 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
1  3  10.10.1.1 (10.10.1.1) 12.7 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
2  1  10.10.4.2 (10.10.4.2) 11.6 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
        entry 2: MPLS Label = 262139, Exp = 7, TTL = 1, S = 1
2  2  10.10.4.2 (10.10.4.2) 13.5 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
        entry 2: MPLS Label = 262139, Exp = 7, TTL = 1, S = 1
2  3  10.10.4.2 (10.10.4.2) 11.9 ms

```




```

        returned MPLS Label Stack Object
          entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262139, Exp = 7, TTL = 1, S = 1
3  1  10.10.6.4 (10.10.6.4) 9.21 ms
        returned MPLS Label Stack Object
          entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
3  2  10.10.6.4 (10.10.6.4) 9.58 ms
        returned MPLS Label Stack Object
          entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
3  3  10.10.6.4 (10.10.6.4) 9.38 ms
        returned MPLS Label Stack Object
          entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
4  1  10.10.10.5 (10.10.10.5) 12.2 ms
        returned MPLS Label Stack Object
          entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
4  2  10.10.10.5 (10.10.10.5) 11.5 ms
        returned MPLS Label Stack Object
          entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
4  3  10.10.10.5 (10.10.10.5) 11.5 ms
        returned MPLS Label Stack Object
          entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262138, Exp = 7, TTL = 1, S = 1
5  1  10.20.1.6 (10.20.1.6) 11.9 ms
5  2  10.20.1.6 (10.20.1.6) 12.2 ms
5  3  10.20.1.6 (10.20.1.6) 13.7 ms

```

3.11.2 VPRN Inter-AS Option B



```

# 12.0R4 default behavior (vc-only)
*A:Dut-A# traceroute 3.3.3.4 source 3.3.4.2 wait 100 no-dns
detail traceroute to 3.3.3.4 from 3.3.4.2, 30 hops max, 40 byte packets
1  1  3.3.4.1  1.97 ms
1  2  3.3.4.1  1.74 ms
1  3  3.3.4.1  1.71 ms
2  1  *
2  2  *
2  3  *
3  1  *
3  2  *
3  3  *
4  1  3.3.3.6  6.76 ms
4  2  3.3.3.6  7.37 ms

```

```

4 3 3.3.3.6 8.36 ms
5 1 3.3.3.4 11.1 ms
5 2 3.3.3.4 9.46 ms
5 3 3.3.3.4 8.28 ms

# Configure icmp-tunneling on C, D, E and F

*A:Dut-A# traceroute 3.3.3.4 source 3.3.4.2 wait 100 no-dns
detail traceroute to 3.3.3.4 from 3.3.4.2, 30 hops max, 40 byte packets
 1 1 3.3.4.1 1.95 ms
 1 2 3.3.4.1 1.85 ms
 1 3 3.3.4.1 1.62 ms
 2 1 10.0.7.3 6.76 ms
    returned MPLS Label Stack Object
      entry 1: MPLS Label = 262143, Exp = 0, TTL = 255, S = 0
      entry 2: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
 2 2 10.0.7.3 6.92 ms
    returned MPLS Label Stack Object
      entry 1: MPLS Label = 262143, Exp = 0, TTL = 255, S = 0
      entry 2: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
 2 3 10.0.7.3 7.58 ms
    returned MPLS Label Stack Object
      entry 1: MPLS Label = 262143, Exp = 0, TTL = 255, S = 0
      entry 2: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
 3 1 10.0.5.4 6.92 ms
    returned MPLS Label Stack Object
      entry 1: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
 3 2 10.0.5.4 7.03 ms
    returned MPLS Label Stack Object
      entry 1: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
 3 3 10.0.5.4 8.66 ms
    returned MPLS Label Stack Object
      entry 1: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
 4 1 3.3.3.6 6.67 ms
 4 2 3.3.3.6 6.75 ms
 4 3 3.3.3.6 6.96 ms
 5 1 3.3.3.4 8.32 ms
 5 2 3.3.3.4 11.6 ms
 5 3 3.3.3.4 8.45 ms

# With ttl-propagate vprn-transit none on PE1 *A:Dut-C# configure router ttl-
propagate vprn-transit none *A:Dut-B# traceroute 3.3.3.4 source 3.3.4.2 wait 100 no-
dns detail traceroute to 3.3.3.4 from 3.3.4.2, 30 hops max, 40 byte packets
 1 1 3.3.4.1 1.76 ms
 1 2 3.3.4.1 1.75 ms
 1 3 3.3.4.1 1.76 ms
 2 1 3.3.3.6 6.50 ms
 2 2 3.3.3.6 6.70 ms
 2 3 3.3.3.6 6.36 ms
 3 1 3.3.3.4 8.34 ms
 3 2 3.3.3.4 7.64 ms
 3 3 3.3.3.4 8.73 ms

# With ttl-propagate vprn-transit all on PE1 *A:Dut-C# configure router ttl-
propagate vprn-transit all *A:Dut-B# traceroute 3.3.3.4 source 3.3.4.2 wait 100 no-
dns detail traceroute to 3.3.3.4 from 3.3.4.2, 30 hops max, 40 byte packets


```

```

1  1  3.3.4.1  1.97 ms
1  2  3.3.4.1  1.77 ms
1  3  3.3.4.1  2.37 ms
2  1  10.0.7.3  9.27 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
2  2  10.0.7.3  6.39 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
2  3  10.0.7.3  6.19 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
3  1  10.0.5.4  6.80 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
3  2  10.0.5.4  6.71 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
3  3  10.0.5.4  6.58 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262140, Exp = 0, TTL = 1, S = 1
4  1  3.3.3.6  6.47 ms
4  2  3.3.3.6  6.75 ms
4  3  3.3.3.6  9.06 ms
5  1  3.3.3.4  7.99 ms
5  2  3.3.3.4  9.31 ms
5  3  3.3.3.4  8.13 ms

```

3.11.3 VPRN Inter-AS Option C and ASBR/ABR/Data Path RR for BGP IPv4 Label Route



12.0R4 default behavior

```

*A:Dut-B# traceroute 16.1.1.1 source 26.1.1.2 detail no-dns
wait 100 traceroute to 16.1.1.1 from 26.1.1.2, 30 hops max, 40 byte packets
1  1  26.1.1.1  1.90 ms
1  2  26.1.1.1  1.81 ms
1  3  26.1.1.1  2.01 ms
2  1  16.1.1.1  6.11 ms
2  2  16.1.1.1  8.35 ms

```

```

2 3 16.1.1.1 5.33 ms

*A:Dut-C# traceroute router 600 26.1.1.2 source 16.1.1.1 detail no-dns
wait 100 traceroute to 26.1.1.2 from 16.1.1.1, 30 hops max, 40 byte packets
1 1 26.1.1.1 5.03 ms
1 2 26.1.1.1 4.60 ms
1 3 26.1.1.1 4.60 ms
2 1 26.1.1.2 6.54 ms
2 2 26.1.1.2 5.99 ms
2 3 26.1.1.2 5.74 ms

# With ttl-propagate vprn-transit all and icmp-tunneling

*A:Dut-B# traceroute 16.1.1.1 source 26.1.1.2 detail no-dns
wait 100 traceroute to 16.1.1.1 from 26.1.1.2, 30 hops max, 40 byte packets
1 1 26.1.1.1 2.05 ms
1 2 26.1.1.1 1.87 ms
1 3 26.1.1.1 1.85 ms
2 1 10.10.4.4 8.42 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262137, Exp = 0, TTL = 1, S = 0
        entry 3: MPLS Label = 262142, Exp = 0, TTL = 1, S = 1
2 2 10.10.4.4 5.85 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262137, Exp = 0, TTL = 1, S = 0
        entry 3: MPLS Label = 262142, Exp = 0, TTL = 1, S = 1
2 3 10.10.4.4 5.75 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262137, Exp = 0, TTL = 1, S = 0
        entry 3: MPLS Label = 262142, Exp = 0, TTL = 1, S = 1
3 1 10.10.1.2 5.54 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262137, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262142, Exp = 0, TTL = 2, S = 1
3 2 10.10.1.2 7.89 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262137, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262142, Exp = 0, TTL = 2, S = 1
3 3 10.10.1.2 5.56 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262137, Exp = 0, TTL = 1, S = 0
        entry 2: MPLS Label = 262142, Exp = 0, TTL = 2, S = 1
4 1 16.1.1.1 9.50 ms
4 2 16.1.1.1 5.91 ms
4 3 16.1.1.1 5.85 ms

# With ttl-propagate vprn-local all
*A:Dut-C# traceroute router 600 26.1.1.2 source 16.1.1.1 detail no-dns
wait 100 traceroute to 26.1.1.2 from 16.1.1.1, 30 hops max, 40 byte packets
1 1 10.10.4.2 4.78 ms
    returned MPLS Label Stack Object
        entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
        entry 2: MPLS Label = 262136, Exp = 7, TTL = 1, S = 0
        entry 3: MPLS Label = 262142, Exp = 7, TTL = 1, S = 1

```

```

1  2  10.10.4.2  4.56 ms
      returned MPLS Label Stack Object
          entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262136, Exp = 7, TTL = 1, S = 0
          entry 3: MPLS Label = 262142, Exp = 7, TTL = 1, S = 1
1  3  10.10.4.2  4.59 ms
      returned MPLS Label Stack Object
          entry 1: MPLS Label = 262143, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262136, Exp = 7, TTL = 1, S = 0
          entry 3: MPLS Label = 262142, Exp = 7, TTL = 1, S = 1
2  1  10.10.6.4  4.55 ms
      returned MPLS Label Stack Object
          entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262142, Exp = 7, TTL = 2, S = 1
2  2  10.10.6.4  4.47 ms
      returned MPLS Label Stack Object
          entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262142, Exp = 7, TTL = 2, S = 1
2  3  10.10.6.4  4.20 ms
      returned MPLS Label Stack Object
          entry 1: MPLS Label = 262138, Exp = 7, TTL = 1, S = 0
          entry 2: MPLS Label = 262142, Exp = 7, TTL = 2, S = 1
3  1  26.1.1.1  4.62 ms
3  2  26.1.1.1  4.41 ms
3  3  26.1.1.1  4.64 ms
4  1  26.1.1.2  5.74 ms
4  2  26.1.1.2  6.22 ms
4  3  26.1.1.2  5.77 ms

```


3.12 Diagnostics Command Reference

- [OAM Commands](#)
- [SAA Commands](#)
- [OAM Performance Monitoring and Binning Commands](#)
- [IP Performance Monitoring Commands](#)
- [Show Commands](#)
- [Clear Commands](#)
- [Monitor Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)

3.12.1 Command Hierarchies

3.12.1.1 OAM Commands

3.12.1.1.1 Base Operational Commands

GLOBAL

- **ping** {*ip-address* | *dns-name*} [{**rapid** | **detail**}] [*tll time-to-live*] [*tos type-of-service*] [*size bytes*] [*pattern pattern*] [*source {ip-address | ipv6-address | dns-name}*] [*interval seconds*] [{**next-hop** *ip-address* | **interface** *interface-name* | **bypass-routing**}] [*count requests*] [**do-not-fragment**] [*router router-instance*] [*time-out timeout*]
- **traceroute** {*ip-address* | *dns-name*} [*tll ttl*] [*wait milli-seconds*] [**no-dns**] [*source src-ip-address*] [*tos type-of-service*] [**router** [*router-instance*]]
- **oam**
 - **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**}]
 - **saa** *test-name* [**owner test-owner**] {**start** | **stop**} [**no-accounting**]

3.12.1.1.2 ATM Diagnostics

GLOBAL

- **oam**
 - **atm-ping** *port-id:vpilvci* [{**end-to-end** | **segment**}] [*dest destination-id*] [*send-count sendcount*] [*timeout seconds*] [*interval seconds*]

3.12.1.1.3 IGMP Snooping

GLOBAL

- **oam**
 - **mfib-ping** **service** *service-id* **source** *src-ip* **destination** *mcast-address* [*size size*] [*tll vc-label-ttl*] [**return-control**] [*interval interval*] [**send-count** *send-count*] [*timeout timeout*]

3.12.1.1.4 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* [profile {in|out}]
 - 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
 - mpls-time-stamp-format {rfc4379 | unix}

3.12.1.1.5 LSP Diagnostics

GLOBAL

- oam
 - lsp-ping *lsp-name* [path *path-name*]
 - 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*]]
 - 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 *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 | non-ip | none}] [path-type active | working | protect]
 - p2mp-lsp-ping {*lsp-name* [p2mp-instance *instance-name* [s2l-dest-address *ip-address* [...(up to 5 max)]]] [ttl *label-ttl*] [fc *fc-name* [profile {in | out}]] [size *octets*] [timeout *timeout*] [detail]}

- **p2mp-lsp-ping** **ldp** *p2mp-identifier* [**vpn-recursive-fec**] [**sender-addr** *ip-address*] [**leaf-addr** *ip-address* [... up to 5 max]] [**fc** *fc-name* [**profile** {**in** | **out**}}] [**size** *octets*] [**timeout** *timeout*] [**detail**]
- **p2mp-lsp-ping** **ldp-ssm** **source** {*ip-address* | *ipv6-address*} **group** {*mcast-address* | *mcastv6-address*} [**router** {*router-instance* | **service-name** *service-name*}] [**sender-addr** *ip-address*] [**leaf-addr** *ip-address* [...up-to-5 max]]] [**fc** *fc-name* [**profile** {**in** | **out**}}] [**size** *octets*] [**timeout** *timeout*] [**detail**]
- **p2mp-lsp-trace** *lsp-name* **p2mp-instance** *instance-name* **s2l-dest-address** *ip-address* [**fc** *fc-name* [**profile** {**in** | **out**}}] [**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**]

3.12.1.1.6 SDP Diagnostics

GLOBAL

— oam

- **sdp-mtu** *orig-sdp-id* **size-inc** *start-octets* *end-octets* [**step** *step-size*] [**timeout** *seconds*] [**interval** *seconds*]
- **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*]

3.12.1.1.7 Common Service Diagnostics

GLOBAL

— oam

- **ancp** {**subscriber** *sub-ident-string* | **ancp-string** *ancp-string*} **loopback** [**count** *count*] [**timeout** *seconds*] [**alarm**]
- **ancp** **subscriber** *sub-ident-string* **loopback** [**send-count** *send-count*] [**timeout** *seconds*] [**alarm**]
- **svc-ping** {*ip-addr* | *dns-name*} **service** *service-id* [**local-sdp**] [**remote-sdp**]
- **host-connectivity-verify** **service** *service-id* [**sap** *sap-id*]
- **host-connectivity-verify** **subscriber** *sub-ident-string* [**sla-profile** *sla-profile-name*]
- **dns** **target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]
- **vprn-ping** *service-id* **source** *src-ip* **destination** *ip-address* [**fc** *fc-name* [**profile** {**in** | **out**}}] [**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**send-count** *send-count*] [**timeout** *timeout*]
- **vprn-trace** *service-id* **source** *src-ip* **destination** *ip-address* [**fc** *fc-name* [**profile** {**in** | **out**}}] [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**return-control**] [**probe-count** *send-count*] [**interval** *seconds*] [**timeout** *timeout*]

3.12.1.1.8 VLL Diagnostics

GLOBAL

— oam

- **vccv-ping** *sdp-id:vc-id* [**reply-mode** *ip-routed* | **control-channel**][**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*][**target-fec-type** *static-pw-fec* **agi** *agi-value* **pw-path-id-saii** *src-global-id:src-node-id:src-ac-id* **pw-path-id-taii** *dest-global-id:dest-node-id:dest-ac-id*] [**count** *send-count*] [**fc** *fc-name*] [**profile** *in|out*]] [**interval** *interval*] [**size** *octets*] [**timeout** *timeout*] [**ttl** *vc-label-ttl*]
- **vccv-ping** **static** *sdp-id:vc-id* [**target-fec-type** *pw-id* **sender-src-address** *ip-address* **remote-dst-address** *ip-address* **pw-id** *value* **pw-type** *value*] [**dest-global-id** *global-id* **dest-node-id** *node-id*] [**assoc-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-ping** **spoke-sdp-fec** *spoke-sdp-fec-id* [**saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id*] [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr*] [**reply-mode** *{ip-routed | control-channel}*] [**fc** *fc-name*] [**profile** *{in | out}*]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**ttl** *vc-label-ttl*]
- **vccv-ping** **saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr*] [**reply-mode** *{ip-routed | control-channel}*] [**fc** *fc-name*] [**profile** *{in | out}*]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**ttl** *vc-label-ttl*]
- **vccv-trace** *sdp-id:vc-id* [**reply-mode** *ip-routed* | **control-channel**] [**target-fec-type** *static-pw-fec* **agi** *agi-value* **pw-path-id-saii** *src-global-id:src-node-id:src-ac-id* **pw-path-id-taii-type2** *dest-global-id:dest-node-id:dest-ac-id*] [**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*]
- **vccv-trace** **static** *sdp-id:vc-id* [**assoc-channel** *ipv4* | **non-ip**] [**src-ip-address** *ipv4-address*] [**target-fec-type** *pw-id* **sender-src-address** *ip-address* **remote-dst-address** *ip-address* **pw-id** *value* **pw-type** *value*] [**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*]
- **vccv-trace** **spoke-sdp-fec** *spoke-sdp-fec-id* [**saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id*]] [**size** *octets*]] [**min-ttl** *min-vc-label-ttl*]] [**max-ttl** *max-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*] [**profile** *{in | out}*]] [**detail**]
- **vccv-trace** **saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id* [**size** *octets*]] [**min-ttl** *min-vc-label-ttl*]] [**max-ttl** *max-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*] [**profile** *{in | out}*]] [**detail**]

GLOBAL

— oam

- **cpe-ping** **service** *service-id* **destination** *dst-ieee-address* **source** *ip-address* [**source-mac** *ieee-address*]] [**ttl** *vc-label-ttl*] [**send-count** *send-count*] [**return-control**] [**interval** *interval*]
- **mac-ping** **service** *service-id* **destination** *dst-ieee-address* [**source** *src-ieee-address*] [**fc** *fc-name*] [**profile** *in | out*]] [**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*] [**profile** *in* | *out*] [**size** *octets*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**probe-count** *send-count*] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]

GLOBAL

- oam
 - **vxlan-ping** **test-id** *test-id* **service** *vpls-service-id* **dest-vni** *vxlan-network-id* **outer-ip-destination** *ipv4-address* [**outer-ip-source-udp** *udp-port-number*] [**outer-ip-ttl** *time-to-live*][**inner-l2** *ieee-address*] [**inner-ip-source** *ipv4-address*][**inner-ip-destination** *ipv4-address*] [**i-flag-on**] [**end-system** *ieee-address*] [**send-count** *packets*] [**interval** *interval-time*][**timeout** *timeout-time*] [**padding** *tlv-size*] [**reflect-pad**] [**fc** *fc-name*] [**profile** *{in|out}*] [**reply-mode** *{overlay|udp}*]

3.12.1.1.9 Ethernet in the First Mile (EFM) Commands

GLOBAL

- oam
 - **efm** *port-id*
 - **local-loopback** *{start | stop}*
 - **remote-loopback** *{start | stop}*

3.12.1.1.10 ETH-CFM OAM Commands

configure

- eth-cfm
 - domain
 - association
 - bridge
 - **id-permission** *chassis*
 - **no id-permission**
 - **facility-id-permission** *chassis*
 - **no facility-id-permission**
 - system
 - **sender-id***[system | local* *local-name**]*
 - **no sender-id**

oam

- eth-cfm
 - **eth-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*] [**data-length** *data-length*]
 - **linktrace** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**tll** *tll-value*]
 - **loopback** *{mac-address|multicast}* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**send-count** *send-count*] [**size** *data-size*] [**priority** *priority*]
 - **one-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]

- **two-way-delay-test** *mac-address mep mep-id domain md-index association ma-index* [**priority** *priority*]
- **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*]

3.12.1.2 SAA Commands

GLOBAL

- **oam**
- [**no**] **saa** *test-name* [**owner** *test-owner*] (**start** | **stop**) [**no-accounting**]

configure

- **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** [**auto** | **drop** | **keep**]
 - [**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*] [**profile** [**in** | **out**]] [**ttl** *vc-label-ttl*] [**send-count** *send-count*] [**send-control**] [**return-control**] [**time-out** *interval*] [**interval** *interval*]
 - **dns** **target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**send-count** *send-count*] [**time-out** *timeout*] [**interval** *interval*]
 - **eth-cfm-linktrace** *mac-address mep mep-id domain md-index association ma-index* [**ttl** *tll-value*] [**fc** *{fc-name}*] [**profile** [**in**|**out**]] [**send-count** *send-count*] [**time-out** *interval*] [**interval** *interval*]
 - **eth-cfm-loopback** *mac-address mep mep-id domain md-index association ma-index* [**size** *data-size*] [**fc** *{fc-name}*] [**profile** [**in**|**out**]] [**send-count** *send-count*][**time-out** *interval*] [**interval** *interval*]

- **eth-cfm-two-way-delay** *mac-address mep mep-id domain md-index association ma-index* [*fc {fc-name}*] [*send-count send-count*][*time-out interval*] [*interval interval*]
- **eth-cfm-two-way-slm** *mac-address mep mep-id domain md-index association ma-index* [*fc {fc-name}*] [*send-count send-count*][*size data-size*][*time-out timeout*] [*interval interval*]
- **icmp-ping** *mac-address mep mep-id domain md-index association ma-index* [*fc {fc-name}*] [*profile {in|out}*]] [*send-count send-count*] [*time-out interval*] [*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 | dns-name*] [*interval interval*] [*{next-hop ip-address} | {interface interface-name} | bypass-routing*] [*count requests*] [*do-not-fragment*] [*router router-instance*] [*time-out interval*]
- **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*]
- **lsp-ping** *lsp-name* [*path path-name*]
- **lsp-ping static** *lsp-name* [*dest-global-id global-id dest-node-id node-id*] [*control-channel none | non-ip*] [*path-type active | working | protect*] [*fc fc-name*] [*profile in | out*] [*interval interval*] [*send-count send-count*] [*size octets*] [*src-ip-address ip-address*] [*timeout timeout*] [*ttl label-ttl*] [*detail*]
- **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*]]
- **lsp-ping static** *lsp-name* [*assoc-channel none|non-ip*] [*dest-global-id global-id dest-node-id node-id*] [*path-type active | working | protect*]
- **lsp-trace** *lsp-name* [*path path-name*]
- **lsp-trace static** *lsp-name* [*control-channel non | non-ip*] [*force*] [*path-type active | working | protect*] [*detail*] [*fc fc-name*] [*profile in|out*]] [*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*] [*downstream-map-tlv dsmmap | ddmap*] [*detail*]
- **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 none|non-ip*] [*path-type active | working | protect*]
- **mac-ping service** *service-id destination ieee-address* [*source src-ieee-address*] [*fc fc-name*] [*profile {in | out}*]] [*size octets*] [*ttl vc-label-ttl*] [*send-count send-count*] [*send-control*] [*return-control*] [*interval interval*] [*time-out interval*]
- **mac-trace service** *service-id destination ieee-address* [*source src-ieee-address*] [*fc fc-name*] [*profile {in | out}*]] [*size octets*] [*min-ttl min-label-ttl*] [*max-ttl max-label-ttl*] [*probe-count send-count*] [*send-control*] [*return-control*] [*interval interval*] [*time-out timeout*]

- **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name*] [**profile** {*in* | *out*}] [**size** *octets*] [**send-count** *send-count*] [**time-out** *interval*] [*interval interval*]
- **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*] [**profile** {*in* | *out*}] [**size** *octets*] [**send-count** *send-count*] [**time-out** *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**] [**time-out** *timeout-value*] [*interval interval-value*] [**fc** *fc-name*] [**profile** {*in* | *out*}] [**detail**]
- **vprn-ping** *service-id* **source** *src-ip* **destination** *dst-ip* [**fc** *fc-name*] [**profile** {*in* | *out*}] [**size** *size*] [**ttl** *vc-label-ttl*] [**send-count** *send-count*] [**return-control**] [**time-out** *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*] [**probe-count** *send-count*] [**return-control**] [**time-out** *timeout*] [*interval interval*]

3.12.1.3 OAM Performance Monitoring and Binning Commands

GLOBAL

- **oam**
- **oam-pm** *session session-name* {*dmm* | *lmm* | *slm* | *twamp-light*} {*start* | *stop*}

configure

- **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**]
- **bin-type** {*fd* | *fdr* | *ifdv*}
- **bin** *bin-number* **lower-bound** *microseconds*
- **delay-event** {*forward* | *backward* | *round-trip*} **lowest-bin** *bin-number* **threshold** *raise-threshold* [**clear** *clear-threshold*]
- **no delay-event** {*forward* | *backward* | *round-trip*}
- **delay-event-exclusion** {*forward* | *backward* | *round-trip*} **lowest-bin** *bin-number*
- **no delay-event-exclusion** {*forward* | *backward* | *round-trip*}
- **exclude-from-avg** {*forward* | *backward* | *round-trip*} **bins** *bin-numbers*
- **no exclude-from-avg** {*forward* | *backward* | *round-trip*}
- [**no**] **description** *description-string*
- [**no**] **shutdown**
- **session** *session-name* **test-family** *ethernet* [**session-type** {*proactive* | *on-demand*}] **create**
- **no session** *session-name*
- **bin-group** *bin-group-name*
- **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**
 - **lmm** [test-id *test-id*] create
 - **no lmm**
 - **availability**
 - **flr-threshold** *percentage*
 - **no flr-threshold**
 - **shutdown**
 - **no shutdown**
 - **timing** *frames-per-delta-t frames consec-delta-t deltas chli-threshold threshold*
 - **no timing**
 - **interval** *milliseconds*
 - **no interval**
 - **loss-events**
 - **avg-flr-event** {forward | backward} threshold *raise-threshold-percent* [clear *clear-threshold-percent*]
 - [no] **avg-flr-event** {forward | backward}
 - **chli-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **chli-event** {forward | backward | aggregate}
 - **hli-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **hli-event** {forward | backward | aggregate}
 - **unavailability-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **unavailability-event** {forward | backward | aggregate}
 - **undet-availability-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **undet-availability-event** {forward | backward | aggregate}
 - **undet-unavailability-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **undet-unavailability-event** {forward | backward | aggregate}
 - [no] **shutdown**
 - **test-duration** *seconds*
 - **no test-duration**
 - **priority** *priority*
 - **no priority**
 - **slm** [test-id *test-id*] create
 - **no slm**
 - **data-tlv-size** *octets*

- **no data-tlv-size**
- **flr-threshold** *percentage*
- **no flr-threshold**
- **loss-events**
 - **avg-flr-event** {forward | backward} threshold *raise-threshold-percent* [clear *clear-threshold-percent*]
 - [no] **avg-flr-event** {forward | backward}
 - **chli-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **chli-event** {forward | backward | aggregate}
 - [no] **flr-threshold** *percentage*
 - **hli-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **hli-event** {forward | backward | aggregate}
 - **unavailability-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **unavailability-event** {forward | backward | aggregate}
 - **undet-availability-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **undet-availability-event** {forward | backward | aggregate}
 - **undet-unavailability-event** {forward | backward | aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **undet-unavailability-event** {forward | backward | aggregate}
- [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*
- **no meas-interval** {5-mins | 15-mins | 1-hour | 1-day}
 - **accounting-policy** *account-policy-id*
 - **no accounting-policy**
 - **boundary-type** {clock-aligned | test-relative}
 - **no boundary-type**
 - **clock-offset** *seconds*
 - **no clock-offset**
 - **event-mon**
 - **delay-events**
 - [no] **delay-events**
 - **loss-events**
 - [no] **loss-events**
 - [no] **shutdown**
 - **intervals-stored** *intervals-stored intervals*
 - **no intervals-stored**

3.12.1.4 IP Performance Monitoring Commands

3.12.1.4.1 TWAMP

```

configure
  — test-oam
    — twamp
      — server
        — [no] prefix {address/prefix-length} [create]
          — description text
          — no description
          — max-conn-prefix count
          — no max-conn-prefix
          — max-sess-prefix count
          — no max-sess-prefix
          — [no] shutdown
        — inactivity-timeout seconds
        — no inactivity-timeout
        — max-conn-server count
        — no max-conn-server
        — max-sess-server count
        — no max-sess-server
        — shutdown

```

3.12.1.4.2 TWAMP Light

```

configure
  — router
    — twamp-light
      — reflector [udp-port udp-port-number] [create]
      — no reflector
        — description description
        — no description
        — prefix {ip-prefix/prefix-length} [create]
        — no prefix
          — description description
          — no description
        — [no] shutdown

configure
  — service
    — vprn
      — [no] twamp-light
        — reflector [udp-port udp-port-number] [create]
        — no reflector
          — description description
          — no description
          — prefix {ip-prefix/prefix-length} [create]

```

```

— no prefix
— description description
— no description
— [no] shutdown

configure
— test-oam
— twamp
— twamp-light
— inactivity-timeout seconds
— no inactivity-timeout

configure
— oam-pm
— session
— ip
— destination ip-address
— no destination
— dest-udp-port udp-port-number
— no dest-udp-port
— fc fc-name
— no fc
— forwarding {next-hop ip-address | interface interface-name | bypass-
routing}
— no forwarding
— profile {in | out}
— no profile
— router {base | routing-instance | 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
— loss
— flr-threshold percentage
— [no] flr-threshold
— timing frames-per-delta-t frames consec-delta-t deltas
chli-threshold threshold
— [no] timing
— loss-events
— avg-flr-event {forward | backward} threshold raise-
threshold-percent [clear clear-threshold-percent]
— [no] avg-flr-event {forward | backward}
— chli-event {forward|backward|aggregate} threshold
raise-threshold [clear clear-threshold]
— [no] chli-event {forward|backward|aggregate}
— hli-event {forward|backward|aggregate} threshold
raise-threshold [clear clear-threshold]

```

-
- [no] **hli-event** {forward|backward|aggregate}
 - **unavailability-event** {forward|backward|aggregate}
threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **unavailability-event**
{forward|backward|aggregate}
 - **undet-availability-event**
{forward|backward|aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **undet-availability-event**
{forward|backward|aggregate}
 - **undet-unavailability-event**
{forward|backward|aggregate} threshold *raise-threshold* [clear *clear-threshold*]
 - [no] **undet-unavailability-event**
{forward|backward|aggregate}
 - **pad-size** *octets*
 - **no pad-size**
 - **pad-size**
 - **record-stats** {delay | loss | delay-and-loss}
 - [no] **record-stats**
 - [no] **shutdown**
 - **test-duration** *seconds*
 - **no test-duration**

3.12.1.5 Show Commands

show

- **eth-cfm**
 - **association** [*ma-index*] [**detail**]
 - **cfm-stack-table**
 - **cfm-stack-table** port [{**all-ports** | **all-sdps** | **all-virtuals**}] [**level 0..7**] [**direction up** | **down**]
 - **cfm-stack-table** port-id [**vlan** *qtag*] [**level 0..7**] [**direction up** | **down**]
 - **cfm-stack-table** sdp *sdp-id[:vc-id]* [**level 0..7**] [**direction up** | **down**]
 - **cfm-stack-table** virtual *service-id* [**level 0..7**]
 - **cfm-stack-table** facility [{**all-ports**|**all-lags**|**all-lag-ports**|**all-tunnel-meps**| **all-router-interfaces**}] [**level 0..7**] [**direction up**|**down**]
 - **cfm-stack-table** facility collect-lmm-stats
 - **cfm-stack-table** facility lag *id* [**tunnel 1..4094**] [**level 0..7**] [**direction up**|**down**]
 - **cfm-stack-table** facility port *id* [**level 0..7**] [**direction up**|**down**]
 - **cfm-stack-table** facility router-interface *ip-int-name* [**level 0..7**] [**direction up**|**down**]
 - **domain** [*md-index*] [**association** *ma-index* | **all-associations**] **statistics** [**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**
 - **statistics**
 - **system-config**
 - **system-info**
- **saa** [*test-name* [**owner** *test-owner*]]
- **test-oam**
 - **ldp-treetrace** [**prefix** *ip-prefix/mask*] [**detail**]
 - **twamp**
 - **client** {**all** | *ip-address*}
 - **server** {**all** | **prefix** *ip-prefix/prefix-length* | **capability**}
 - **twamp-light**
 - **reflectors**

show

- **oam-pm**
 - **bin-group** [*bin-group-number*] [**detail**]
 - **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** | **detectable-tx-errors**}
 - **statistics** *session session-name* {**dmm**| **lmm** | **slm** | **twamp-light**} **meas-interval** {**raw** | **5mins** | **15-min** | **1-hour** | **1-day**} [{**all** | **bins** | **summary**}] **interval-number** [*interval-number* [{**delay** | **loss**}]

3.12.1.6 Clear Commands

```

clear
  — saa [test-name [owner test-owner]]

clear
  — oam-pm
    — session session-name {dmm | lmm | slm | twamp-light}

clear
  — eth-cfm
    — auto-discovered-meps mep-id domain md-index association
    — statistics

```

3.12.1.7 Monitor Commands

```

monitor
  — oam-pm
    — session session-name {dmm | lmm | slm | twamp-light}

```

3.12.1.8 Debug Commands

```

debug
  — eth-cfm
    — [no] mep mep-id domain md-index association ma-index
      — packet all
      — packet cfm-opcode opcode [opcode ... (up to 5 max.)]
      — no packet
    — [no] mip domain md-index association ma-index
      — packet all
      — packet cfm-opcode opcode [opcode ... (up to 5 max.)]
      — no packet
  — oam
    — lsp-ping-trace [tx | rx | both] [raw | detail]
    — no lsp-ping-trace

```

3.12.1.9 Tools Commands

```

tools
  — dump
    — eth-cfm
      — debug-packet [clear]
      — top-active-mep [rx-sort | tx-sort] [clear]

```

3.12.2 Command Descriptions

3.12.2.1 OAM and SAA Commands

3.12.2.1.1 Generic Commands

shutdown

Syntax	[no] shutdown
Context	config>saa>test
Description	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	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.

3.12.2.1.2 OAM Commands

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	oam
Description	This command performs DNS name resolution. If <code>ipv4-a-record</code> is specified, <code>dns-names</code> are queried for A-records only. If <code>ipv6-aaaa-record</code> is specified, AAAA-records are queried first, and if a successful reply is not received, the <code>dns-server</code> is queried for A-records (applies to the 7750 SR and 7950 XRS).
Parameters	send-count <i>send-count</i> — 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 to 100 <i>ip-address</i> — The IP or IPv6 address of the primary DNS server.

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 to FFFF]H
 d: [0 to 255]D



Note: `ipv6-address` applies to the 7750 SR and 7950 XRS only.

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 to 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 to 10

record-type — Specifies a record type (applies to the 7750 SR and 7950 XRS only).

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 { <i>ip-address</i> <i>dns-name</i> } [{ rapid detail }] [ttl <i>time-to-live</i>] [tos <i>type-of-service</i>] [size <i>bytes</i>] [pattern <i>pattern</i>] [source <i>ip-address</i> <i>dns-name</i>] [interval <i>seconds</i>] [{ next-hop <i>ip-address</i> interface <i>interface-name</i> bypass-routing }] [count <i>requests</i>] [do-not-fragment] [router <i>router-instance</i>] [timeout <i>timeout</i>]
Context	<GLOBAL>
Description	This command verifies the reachability of a remote host.
Parameters	<i>ip-address</i> — 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[-interface] x:x:x:x:x:x.d.d.d.d[-interface] x: [0 to FFFF]H d: [0 to 255]D interface: 32 characters maximum, mandatory for link local addresses ipv6-address: x:x:x:x:x:x:x x:x:x:x:x:x.d.d.d.d x: [0 to FFFF]H d: [0 to 255]D



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

dns-name — The DNS name of the far-end device to which to send the **svc-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 to 128

tos *type-of-service* — Specifies the service type.

Values 0 to 255

size *bytes* — The request packet size in bytes, expressed as a decimal integer.

Values 0 to 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 to 65535

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
 x:x:x:x:x:x.d.d.d.d
 x: [0 to FFFF]H
 d: [0 to 255]D
ipv6-address: x:x:x:x:x:x:x
 x:x:x:x:x:x.d.d.d.d
 x: [0 to FFFF]H
 d: [0 to 255]D



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

router *router-instance* — Specifies the router name or service ID.

Values *router-name:* Base, management
service-id: 1 to 2147483647

Default Base

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.

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 to FFFF]H
 d: [0 to 255]

Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

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 to 100000

Default 5

do-not-fragment — Sets the DF (Do Not Fragment) bit in the ICMP ping packet (does not apply to ICMPv6).

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 to 10

traceroute

Syntax	traceroute [<i>ip-address</i> [<i>dns-name</i>]] [ttl <i>ttl</i>] [wait <i>milli-seconds</i>] [no-dns] [source <i>ip-address</i>] [tos <i>type-of-service</i>] [router <i>router-instance</i>]
Context	<GLOBAL>
Description	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.

Values

ipv4-address: a.b.c.d
 ipv6-address: x:x:x:x:x:x:x
 x:x:x:x:x:d.d.d.d
 x: [0 to FFFF]H
 d: [0 to 255]



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

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 to 255

wait *milli-seconds* — The time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Default 5000

Values 1 to 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 to 255

router *router-name* — Specifies the alphanumeric character string up to 32 characters.

Default Base

router *service-id* — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR and 7950 XRS router on which this service is defined.

Values 1 to 2147483647

p2mp-lsp-ping

Syntax	p2mp-lsp-ping {<i>lsp-name</i> [p2mp-instance <i>instance-name</i> [s2l-dest-address <i>ip-address</i> [...(up to 5 max)]]] [ttl <i>label-ttl</i>]} [fc <i>fc-name</i> [profile {in out}]] [size <i>octets</i>] [timeout <i>timeout</i>] [detail] p2mp-lsp-ping {ldp <i>p2mp-identifier</i> [vpn-recursive-fec] [sender-addr <i>ip-address</i>] [leaf-addr <i>ip-address</i> [...up to 5 max]]} [fc <i>fc-name</i> [profile {in out}]] [size <i>octets</i>] [timeout <i>timeout</i>] [detail] p2mp-lsp-ping {ldp-ssm source {<i>ip-address</i> <i>ipv6-address</i>} group {<i>mcast-address</i> <i>mcast-v6-address</i>} [router {<i>router-instance</i> service-name <i>service-name</i>}] [sender-addr <i>ip-address</i>] [leaf-addr <i>ip-address</i> [...up-to-5 max]]} [fc <i>fc-name</i> [profile {in out}]] [size <i>octets</i>] [timeout <i>timeout</i>] [detail]
Context	oam
Description	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. LDP P2MP generic-identifier along with source IP address of the head-end node can be used to uniquely identify LDP P2MP LSP in a 7750 SR or 7950 XRS network. LDP p2mp-identifier is a mandatory parameter to test LSP ping. LDP P2MP identifier specified to configure a tunnel-interface on head-end node must be used as p2mp-identifier to test a particular LSP. 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 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. Note that p2mp-lsp-ping is not supported in a VPLS/B-VPLS PMSI context. 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	fc <i>fc-name</i> — The fc and profile parameters are used to indicate the forwarding class and profile 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 and profile 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 and profile 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 11](#) summarizes this behavior.

Table 11 p2mp-lsp-ping Request Packet and Behavior

cpm (sender node)	echo request packet: • packet{tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
outgoing interface (sender node)	echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	echo request packet: • packet{tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
cpm (responder node)	echo reply packet: • packet{tos=1, fc2, profile2}
outgoing interface (responder node)	echo reply packet: • pkt queued as {fc2, profile2} • 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)	echo reply packet: • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface

Default be

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

ldp p2mp-identifier — Identifier to specify a LDP P2MP LSP to ping (applies to the 7750 SR and 7950 XRS only).

Values The p2mp-identifier must be a 32 bit integer.

leaf-addr *ip-address* [*ip-address up to 5 max*] — Specifies the list of egress LER system addresses which are required to reply to LSP ping echo request message (applies to the 7750 SR and 7950 XRS only).

Values ipv4-address: a.b.c.d

lsp-name — Name that identifies an P2MP LSP to ping. The LSP name can be up to 32 characters long.

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.

profile {in | out} — The profile of the LSP ping echo request message.

s2l-dest-addr *ip-address* [*ip-address...up to 5*] — Specifies the list of egress LER system addresses which are required to reply to the LSP ping echo request message.

Default out

sender-addr *ip-address* — Specifies any local IP sender-addr for mLDP (applies to the 7750 SR and 7950 XRS only).

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.

Default 1 octet

Values 1 to 9198

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.

Default 10 seconds

Values 1 to 120

ttl *label-ttl* — The TTL value for the MPLS label, expressed as a decimal integer.

Default 255

Values 1 to 255

vpn-recursive-fec — Adds a VPN recursive FEC element to the launched packet (useful for pinging a VPN BGP inter-AS Option B leaf). This parameter issues an OAM **p2mp-lsp-ping** with RFC 6512 VPN recursive opaque FEC type 8.

Refer to the “OAM” subsection of the LDP chapter in the *MPLS Guide* for more information.

ldp-ssm — Configures a specific multicast stream to be tested when using dynamic multicast in mLDP. The source and group addresses correspond to the <S,G> being advertised by this mLDP FEC.

Values		
source	ipv4-address	a.b.c.d
	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x - [0 to FFFF]H d - [0t o 255]D
group	mcast-address	
	mcast-v6-address	
router	router-name	Base management Default - Base
	service-id	[1 to 2147483647]
	service-name	[64 chars max]
sender-addr	ipv4-address	a.b.c.d
leaf-addr	ipv4-address	a.b.c.d

p2mp-lsp-trace

Syntax	p2mp-lsp-trace <i>lsp-name</i> p2mp-instance <i>instance-name</i> s2l-dest-address <i>ip-address</i> [fc <i>fc-name</i> [profile { in out }]] [size <i>octets</i>] [max-fail <i>no-response-count</i>] [probe-count <i>probes-per-hop</i>] [min-ttl <i>min-label-ttl</i>] [max-ttl <i>max-label-ttl</i>] [timeout <i>timeout</i>] [interval <i>interval</i>] [detail]
Context	oam
Description	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>". Note that p2mp-lsp-trace is not supported in a VPLS/B-VPLS PMSI context.

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 **fc** *fc-name* — The fc and profile parameters are used to indicate the forwarding class and profile 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 and profile 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 and profile 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 12](#) summarizes this behavior.

Table 12 2mp-lsp-trace Request Packet and Behavior

cpm (sender node)	echo request packet: • packet{tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
outgoing interface (sender node)	echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	echo request packet: • packet{tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
cpm (responder node)	echo reply packet: • packet{tos=1, fc2, profile2}
outgoing interface (responder node)	echo reply packet: • pkt queued as {fc2, profile2} • 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)	echo reply packet: • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface

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 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.

Default 1

Values 1 to 10

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.

Default 5

Values 1 to 255

max-ttl *max-label-ttl* — Specifies the maximum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 30

Values 1 to 255

min-ttl *min-label-ttl* — Specifies the minimum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 1

Values 1 to 255

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.

Default 1

Values 1 to 10

profile {in | out} — The profile of the LSP trace echo request message.

Default out

s2l-dest-addr *ip-address* — Specifies the egress LER system address of the S2L sub-LSP path which is being traced.

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.

Default 1 octets

Values 1 to 9198

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.

Default 3 seconds

Values 1 to 60

Output

Sample Output

```
*A:Dut-C# oam p2mp-lsp-trace "p2mp_1" p2mp-instance "1" s2l-dest-address 10.20.1.
10.20.1.4 10.20.1.5 10.20.1.6
*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)

Probe returned multiple responses. Result may be inconsistent.

*A:Dut-C#
```

3.12.2.1.3 ATM Diagnostics

The commands described in this section apply only to the 7750 SR.

atm-ping

Syntax	atm-ping <i>port-id</i> : <i>vpilvci</i> [end-to-end segment] [dest <i>destination-id</i>] [send-count <i>send-count</i>] [timeout <i>timeout</i>] [interval <i>seconds</i>]
Context	oam
Description	This command tests ATM path connectivity and round trip time on an ATM VCC.

Parameters *port-id: vpi/vci* — Specifies the ID of the access port of the target VC. This parameter is required.

Values			
<i>port-id</i>	<i>slot/mda/port [.channel]</i>		
	<i>eth-sat-id</i>	<i>esat-id/slot/port</i>	
		<i>esat</i>	keyword
		<i>id</i>	1 to 20
	<i>pxc-id</i>	<i>pxc-id.sub-port</i>	
		<i>pxc</i>	keyword
		<i>id</i>	1 to 64
		<i>sub-port</i>	a, b
<i>aps-id</i>	<i>group-id</i>		
	<i>aps</i>	keyword	
	<i>group-id</i>	1 to 64	
<i>vpi</i>	0 to 4095 (NNI)		
	0 to 255 (UNI)		
<i>vci</i>	1, 2, 5 to 65535		

end-to-end | segment — Specifies whether the ATM OAM loopback cell is destined to the first segment point in the line direction or the PVCC's connection endpoint.

Default end-to-end

dest *destination-id* — Defines the LLID field in an OAM loopback cell. If set to all 1s, only the connection end (end-to-end ping) or segment end (segment ping) will respond to the ping. If the 'segment' parameter is specified and 'dest' is set to a specific destination, only the destination will respond to the ping.

Values A 16 byte octet string, with each octet separated by a colon, if not specified the value of all 0x11 will be used.

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 to 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 to 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 to 10

3.12.2.1.4 Service Diagnostics

ancp

Syntax	ancp { subscriber <i>sub-ident-string</i> ancp-string <i>ancp-string</i> } loopback [count <i>count</i>] [timeout <i>seconds</i>] [alarm] ancp subscriber <i>sub-ident-string</i> loopback [send-count <i>send-count</i>] [timeout <i>seconds</i>] [alarm]
Context	oam
Description	This command sends an OAM request to the access node. ANCP can be used to send OAM messages to the access node. The access node must be able to accept these messages and will signal such support by the capability negotiations. If the operator attempts to send an OAM command to an access node that does not support such command the operation results in an error.
Parameters	subscriber <i>sub-ident-string</i> — Specifies an existing subscriber-id. The node will use the <i>ancp-string</i> associated with the provided subscriber-id to identify the circuit. ancp-string <i>ancp-string</i> — Specifies an existing ANCP string.

send-count *send-count* — Specifies the number of messages the access node will use to test the circuit. If omitted, the number will be determined by the access node via local policy.

Values 1 to 32

timeout *seconds* — Specifies how long the controlling node will wait for a result.

Values 0 to 300

alarm — Specifies that the CLI the result will be returned to the CLI and a trap will be issued to indicate the test finished. If the flag is used through SNMP the results will be available in the results MIB and after the node sent the trap to indicate the results are ready.

loopback — Sends an OAM loopback test request to the access node

sdp-mtu

Syntax **sdp-mtu** *orig-sdp-id* **size-inc** *start-octets end-octets* [**step** *step-size*] [**timeout** *seconds*] [**interval** *seconds*]

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/GRE encapsulation from the far-end router. OAM request messages sent within an IP/GRE 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/GRE encapsulation from the far-end router.

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/GRE 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 to 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 to 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 to 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 to 512
- 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 to 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 to 10

Output

Sample Output for SDP MTU Path Test

```
*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
3072	.	Success

Maximum Response Size: 3072

```
*A:Dut-A#
```

svc-ping

Syntax **svc-ping** *ip-address* [**service** *service-id*] [**local-sdp**] [**remote-sdp**]

Context oam

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**:

Local and remote service existence

- Local and remote service state
- Local and remote service type correlation
- Local and remote customer association
- Local and remote service-to-SDP bindings and state
- 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.

Table 13 **Svc-ping**

Field	Description	Values
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.	service-id
Local Service Type	The type of service being tested. If <i>service-id</i> does not exist locally, N/A is displayed.	Epipe, Ipipe, Fpipe, Apipe
		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

Table 13 **Svc-ping (Continued)**

Field	Description	Values (Continued)
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.	Epipe, Ipipe, Fpipe, Apipe
		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
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

Table 13 Svc-ping (Continued)

Field	Description	Values (Continued)
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

Table 13 Svc-ping (Continued)

Field	Description	Values (Continued)
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
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

Table 13 Svc-ping (Continued)

Field	Description	Values (Continued)
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
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

Table 13 **Svc-ping (Continued)**

Field	Description	Values (Continued)
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 its egress service label operational, Up is displayed. If the responding router considers its 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
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 its ingress service label operational, Up is displayed. If the originating router considers its 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

Table 13 **Svc-ping (Continued)**

Field	Description	Values (Continued)
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 its ingress service label operational, Up is displayed. If the remote router considers its 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 *ip-address* — The far-end IP address to which to send the **svc-ping** request message in dotted decimal notation.

service *service-id* — The service ID of the service being tested must be indicated with this parameter. The service ID need not exist on the local router to receive a reply message.

Values 1 to 2147483647

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 *sdp-id* 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 *sdp-id* is the expected *responder-id* within the reply received. The *sdp-id* defines the encapsulation of the SDP tunnel encapsulation used to reach the far end; this can be IP/GRE or MPLS. On originator egress, the service-ID must have an associated VC-Label to reach the far-end 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 GRE encapsulation with the OAM label.

[Table 14](#) indicates whether a message is sent and how the message is encapsulated based on the state of the service ID.

Table 14 Message Encapsulation

Local Service State	local-sdp Not Specified		local-sdp Specified	
	Message Sent	Message Encapsulation	Message Sent	Message Encapsulation
Invalid Local Service	Yes	Generic IP/GRE OAM (PLP)	No	None
No Valid SDP-ID Bound	Yes	Generic IP/GRE OAM (PLP)	No	None
SDP-ID Valid But Down	Yes	Generic IP/GRE OAM (PLP)	No	None
SDP-ID Valid and Up, But No Service Label	Yes	Generic IP/GRE OAM (PLP)	No	None
SDP-ID Valid, Up and Egress Service Label	Yes	Generic IP/GRE 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/GRE 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.

If **remote-sdp** is not specified, the **svc-ping** request message is sent with GRE encapsulation with the OAM label.

[Table 15](#) indicates how the message response is encapsulated based on the state of the remote service ID.

Table 15 Message Response Encapsulation

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)

Table 15 **Message Response Encapsulation (Continued)**

Remote Service State	Message Encapsulation (Continued)	
	remote-sdp Not Specified	remote-sdp Specified
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)

Output**Sample Output**

```
*A:router1> svc-ping far-end 10.10.10.10 service 101 local-sdp remote-sdp
Request Result: Sent - Reply Received
```

```
Service-ID: 101
```

```
Err      Basic Info          Local      Remote
---      -
Type:          TLS           TLS
Admin State:   Up            Up
Oper State:    Up            Up
Service-MTU:   1514          1514
Customer ID:   1001          1001
```

```
Err      System IP Interface Info
```

```
---      -
Local Interface Name: "7750 SR-System-IP-Interface (Up to 32 chars)..."
Local IP Interface State:      Up
Local IP Address:              10.10.10.11
IP Address Expected By Remote: 10.10.10.11
Expected Remote IP Address:    10.10.10.10
Actual Remote IP Address:      10.10.10.10
```

```
Err      SDP-ID Info          Local      Remote
---      -
Path Used:      Yes           Yes
SDP-ID:         123           325
Administrative State: Up       Up
Operative State: Up           Up
Binding Admin State: Up       Up
Binding Oper State: Up        Up
Binding VC-ID:  101           101
```

```
Err      Service Label Information  Label      Source      State
---      -
```

—	Local Egress Label:	45	Signaled	Up
—	Remote Expected Ingress:	45	Signaled	Up
—	Remote Egress:	34	Signaled	Up
—	Local Expected Ingress:	34	Signaled	Up

host-connectivity-verify

Syntax	host-connectivity-verify service <i>service-id</i> [sap <i>sap-id</i>] host-connectivity-verify subscriber <i>sub-indent-string</i> [sla-profile <i>sla-profile-name</i>]
Context	oam
Description	This command triggers the host connectivity verification checks and applies only to the 7450 ESS and 7750 SR.
Parameters	service <i>service-id</i> — Specifies the service ID to diagnose or manage. Values 1 to 2147483647 sap <i>sap-id</i> — Specifies the physical port identifier portion of the SAP definition. sub-profile <i>sub-indent-string</i> — Specifies an existing subscriber-id. sla-profile <i>sla-profile-name</i> — Specifies an existing SLA profile name. The SLA profile is configured in the config>subscr-mgmt>sla-profile context.

vprn-ping

Syntax	vprn-ping <i>service-id</i> source <i>ip-address</i> destination <i>ip-address</i> [fc <i>fc-name</i> [profile [in out]] [size <i>size</i>] [ttl <i>vc-label-ttl</i>] [return-control] [interval <i>interval</i>] [send-count <i>send-count</i>] [timeout <i>timeout</i>]
Context	oam config>saa>test>type
Description	This command performs a VPRN ping and applies only to the 7750 SR and 7950 XRS.
Parameters	service <i>service-id</i> — The VPRN service ID to diagnose or manage. Values <i>service-id</i> : 1 to 2147483647 <i>svc-name</i> : 64 characters maximum source <i>ip-address</i> — Specifies an unused IP address in the same network that is associated with the VPRN. Values ipv4-address: 0.0.0.0 to 255.255.255.255 ipv6-address: x:x:x:x:x:x:x x:x:x:x:x:d.d.d.d x: [0 to FFFF]H

d: [0 to 255]D

destination *ip-address* — Specifies the IP address to be used as the destination for performing a VPRN ping operation.

Values 0.0.0.0 to 255.255.255.255

size *octets* — The OAM request packet size in octets, expressed as a decimal integer.

Values 1 to 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 to 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 to 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 to 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 to 100

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

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

profile {in | out} — The profile state of the MPLS echo request encapsulation.

Default out

Output

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* [**profile** {in | out}] [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**return-control**] [**probe-count** *probes-per-hop*] [**interval** *seconds*] [**timeout** *timeout*]

Context oam
 config>saa>test>type

Description This command is used to perform a VPRN trace.

Parameters **service** *service-id* — Specifies the VPRN service ID to diagnose or manage.

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

source *src-ip* — Specifies the IP address for the source IP address in dotted decimal notation.

Values

ipv4-address:	0.0.0.0 to 255.255.255.255
ipv6-address:	x:x:x:x:x:x:x
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

destination *dst-ip* — Specifies the IP address to be used as the destination for performing an vprn-trace operation.

Values 0.0.0.0 to 255.255.255.255

size *octets* — The OAM request packet size in octets, expressed as a decimal integer.

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 to 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 to 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.

probe-count *send-count* — The number of OAM requests sent for a particular TTL value, expressed as a decimal integer.

Default 1

Values 1 to 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 to 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 3

Values 1 to 10

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

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

profile {in | out} — The profile state of the MPLS echo request encapsulation.

Default out

Output

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#
```


3.12.2.1.5 VPLS MAC Diagnostics

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*] [**profile** [**in** | **out**]] [**interval** *interval*] [**send-count** *send-count*]

Context oam
config>saa>test>type

Description 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 to 2147483647
svc-name: 64 characters maximum

destination *ip-address* — Specifies the IP address to be used as the destination for performing an OAM ping operations.

source *ip-address* — Specifies an unused IP address in the same network that is associated with the VPLS or PBB Epipe.

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Values 1 to 255

Default 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. This parameter is only valid for VPLS services.

Default MAC OAM reply sent using the data plane.

source-mac *ieee-address* — Specifies 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 or CFM is used. This parameter is not applicable to CPE ping on Epipes.

fc-name — The forwarding class of the MPLS echo request encapsulation.

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

Default be

profile {in | out} — The profile state of the MPLS echo request encapsulation for VPLS and the ARP packet for PBB Epipe and Epipe VLLs.

Default out

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.

Values 1 to 10

Default 1

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.

Values 1 to 100

Default 1

send-control — This command has been deprecated. SAA test that include this deprecated command will fail the no shutdown.

mac-populate

Syntax **mac-populate** *service-id* **mac** *ieee-address* [**flood**] [**age** *seconds*] [**force**]

Context oam

Description 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 is sent via 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.

Parameters	service <i>service-id</i> — The Service ID of the service to diagnose or manage.
	Values 1 to 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 to 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 sap-id value is not specified the MAC is bound to the CPM or CFM. The originating SHG is 0 (zero). When the target-sap sap-id 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 target ieee-address [flood] [register]</i>
Context	oam
Description	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.

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.

If the register option is provided, the R bit in the Address Delete flags is turned on.

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 is sent via 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 to 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*] [**profile** *in* | *out*] [**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 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 is sent via 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 <i>service-id</i> — The service ID of the service to diagnose or manage. Values 1 to 2147483647 destination <i>ieee-address</i> — The destination MAC address for the OAM MAC request. size <i>octets</i> — 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 to 65535 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 to 255 send-control — This command has been deprecated. SAA test that include this deprecated command will fail the no shutdown.

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

profile {in | out} — The profile state of the MPLS echo request encapsulation.

Default out

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 to 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 to 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 to 10

mac-trace

Syntax	mac-trace service <i>service-id</i> destination <i>ieee-address</i> [size <i>octets</i>] [min-ttl <i>vc-label-ttl</i>] [max-ttl <i>vc-label-ttl</i>] [send-control] [return-control] [source <i>ieee-address</i>] [z-count <i>probes-per-hop</i>] [interval <i>interval</i>] [timeout <i>timeout</i>]
Context	oam config>saa>test>type
Description	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 Nokia 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 via the 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 <i>ieee-address</i> 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 <i>service-id</i> — The Service ID of the service to diagnose or manage. Values 1 to 2147483647 destination <i>ieee-address</i> — The destination MAC address to be traced. size <i>octets</i> — 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 to 65535 min-ttl <i>vc-label-ttl</i> — The minimum TTL value in the VC label for the MAC trace test, expressed as a decimal integer. Default 1 Values 1 to 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 to 255

send-control — This command has been deprecated. SAA test that include this deprecated command will fail the no shutdown.

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 to 100

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 to 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 to 10

vxlan-ping

Syntax **vxlan-ping test-id** *test-id* **service** *vpls-service-id* **dest-vni** *vxlan-network-id* **outer-ip-**

	destination <i>ipv4-address</i> [outer-ip-source-udp <i>udp-port-number</i>] [outer-ip-ttl <i>time-to-live</i>] [inner-l2 <i>ieee-address</i>] [inner-ip-source <i>ipv4-address</i>] [inner-ip-destination <i>ipv4-address</i>] [i-flag-on] [end-system <i>ieee-address</i>] [send-count <i>packets</i>] [interval <i>interval-time</i>] [timeout <i>timeout-time</i>] [padding <i>tlv-size</i>] [reflect-pad] [fc <i>fc-name</i>] [profile { <i>in</i> <i>out</i> }] [reply-mode { <i>overlay</i> <i>udp</i> }]
Context	oam
Description	Operational command used to validate the VXLAN Tunnel Endpoint (VxLAN) connectivity between peers.
Parameters	test-id <i>test-id</i> — Specifies a value to identify the originator handle of the specific request. Each active test requires a unique test identifier. Values 1 to 2147483647 Default must be specified service <i>vpls-service-id</i> — Specifies the VPLS service used to launch the request and by extension pickup the source VNI information. Values 1 to 2147483647 service-name: 64 char max Default must be specified dest-vni <i>vxlan-network-id</i> — Specifies the target Vxlan network identifier on the terminating VTEP. Values 1 to 16777215 Default must be specified outer-ip-destination <i>ipv4-address</i> — Specifies the IPv4 address of the terminating VTEP. Values format a.b.c.d Default must be specified outer-ip-source-udp <i>udp-port-number</i> — Optional Outer source UDP port number. Values 1 to 65535 Default System-generated UDP port number outer-ip-ttl <i>time-to-live</i> — Specifies the optional outer time to live. Values 1 to 255 Default 255 inner-l2 <i>ieee-address</i> — Optional destination mac address used in the inner VxLAN header. Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx Default 00:00:00:00:00:00

inner-ip-source *ipv4-address* — Optional inner source IPv4 address.

Values format a.b.c.d

Default System IPv4 Address

inner-ip-destination *ipv4-address* — Optional inner destination IPv4 address, must be in the range 127/8.

Values In the 127.0.0.0/8 range

Default 127.0.0.1

reply-mode *overlay* | *udp* — Optional keyword that instructs the responder how to route the VxLAN response.

Values *udp*: responds using UDP over the IP network. The default must be changed if the VTEP uses anything other than an IPv4 System Address as the source.

overlay: responds using the VXLAN overlay for the service

Default udp

i-flag-on — Optional keyword to set the VNI Validation bit to 1 indicating the OAMPDU contains a valid VNI.

Default i-flag set to “0” which prevents the OAMPDU from being forwarded beyond the terminating VTEP.

end-system *ieee-address* — Optional command to include the sub TLV to validate an end system MAC address in the FDB. Only one MAC address may be included.

Default No end system TLV included

send-count *packets* — Optional command to adjust the number of VxLAN ping requests transmitted.

Values 1 to 1024

Default 1

interval *interval-time* — Optional command to adjust the probe interval.

Values 0.1 | 1 to 10

Default 1 second

timeout *timeout-time* — Optional command to adjust the default 5 second packet timeout value.

Values 1 to 10

Default 5 seconds

padding *tlv-size* — Optional command to include the Pad TLV. The number of octets that defines the entire size of the pad TLV, including the type (2B), the length field (2B), the padding (variable).

Values 0 | 5 to 2000

Default 0 (not included)

reflect-pad — Optional keyword used to instruct the responder to include the pad-tlv in the echo response. This option is not supported when the reply mode is “UDP”.

Default pad is not reflected

fc *fc-name* — Optional command used to indicate the forwarding class that will be exposed to the QoS policy as input into generating the outer CoS.

Values be | l2 | af | l1 | h2 | ef | h1 | nc

Default be

profile {in | out} — Optional keyword used to define the frame’s disposition that will be exposed to the QoS policy as input into generating the outer CoS.

Default in

3.12.2.1.6 IGMP Snooping Diagnostics

mfib-ping

Syntax	mfib-ping service <i>service-id</i> source <i>src-ip</i> destination <i>mcast-address</i> [size <i>size</i>] [ttl <i>vc-label-ttl</i>] [return-control] [interval <i>interval</i>] [send-count <i>send-count</i>] [timeout <i>timeout</i>]
Context	oam
Description	The mfib-ping utility determines the list of SAPs which egress a certain IP multicast stream (identified by source unicast and destination multicast IP addresses) within a VPLS service. An mfib-ping packet is always sent via the data plane. An mfib-ping is forwarded across the VPLS following the MFIB. If an entry for the specified source unicast and destination multicast IP addresses exist in the MFIB for that VPLS, 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 the specified IP multicast stream. An mfib-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 reply is sent using the data plane.
Special Cases	MFIB 224.0.0.X pings — Mfib-ping requests directed to a destination address in the special 224.0.0.X range are flooded throughout the service flooding domain and will receive a response from all operational SAPs. Note that SAPs that are operationally down do not reply. If EMG is enabled, mfib-ping will return only the first SAP in each chain.
Parameters	service <i>service-id</i> — Specifies the service ID of the VPLS to diagnose or manage. Values 1 to 2147483647 source <i>src-ip</i> — Specifies the source IP address for the OAM request.

destination *mcast-address* — Specifies the destination multicast address for the OAM request.

size *size* — The multicast 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 to 65535

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM request, expressed as a decimal integer.

Default 255

Values 1 to 255

return-control — Specifies the OAM reply has to be sent using the control plane instead of the data plane.

Default OAM reply is sent using the data 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 to 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 to 100

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 next 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 to 100

Output

Sample Output for Multicast FIB Connectivity Test

```
A:ALA-A# oam mfib-ping service 10 source 10.10.10.1 destination 225.0.0.1 count 2
Seq Node-id                                     Path      Size  RTT
-----
[Send request Seq. 1.]
1  51.51.51.51:sap1/1/1                         Self       100   0ms
1  54.54.54.54:sap1/1/2                         In-Band    100  20ms
1  54.54.54.54:sap1/1/3                         In-Band    100  10ms
1  52.52.52.52:sap1/1/3                         In-Band    100  20ms
[Send request Seq. 2.]
2  51.51.51.51:sap1/1/1                         Self       100   0ms
2  52.52.52.52:sap1/1/2                         In-Band    100  10ms
2  54.54.54.54:sap1/1/2                         In-Band    100  10ms
2  52.52.52.52:sap1/1/3                         In-Band    100  20ms
2  54.54.54.54:sap1/1/3                         In-Band    100  30ms
-----
A:ALA-AIM# oam mfib-ping service 1 source 11.11.0.0 destination 224.0.0.1
Seq Node-id                                     Path      Size  RTT
-----
[Send request Seq. 1.]
1  10.20.1.3:sap1/1/5:1                         Not in MFIB Self    40    0ms
1  10.20.1.3:sap1/1/2:1                         Self       40   10ms
[Echo replies received: 2]
-----
A:ALA-AIM#
```

3.12.2.1.7 EFM Commands

efm

Syntax **efm** *port-id*

Context oam>efm

Description 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.

When EFM OAM is disabled or shutdown on a port, the dying gasp flag for the OAMPDU is set for the OAMPDUs sent to the peer. This speeds up the peer loss detection time.

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



Note: On the 7950 XRS, The XMA ID takes the place of the MDA.

<i>port-id</i>	<i>slot/mda/port [.channel]</i>		
<i>eth-sat-id</i>	<i>esat-id/slot/port</i>		
	<i>esat</i>	keyword	
	<i>id</i>		1 to 20
<i>pxc-id</i>	<i>pxc-id.sub-port</i>		
	<i>pxc</i>	keyword	
	<i>id</i>		1 to 64
	<i>sub-port</i>		a, b

local-loopback

Syntax **local-loopback {start | stop}**

Context oam>efm

Description This command enables local loopback tests on the specified port.

remote-loopback

Syntax **remote-loopback {start | stop}**

Context oam>efm

Description 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.

In order for EFM OAM tunneling to function properly, EFM OAM tunneling should be configured for VLL services or a VPLS service with two SAPs only.

3.12.2.1.8 ETH-CFM OAM Commands

linktrace

Syntax	linktrace <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [ttl <i>ttl-value</i>]
Context	oam>eth-cfm
Description	The command specifies to initiate a linktrace test.
Parameters	<i>mac-address</i> — Specifies a unicast destination MAC address. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 ttl <i>ttl-value</i> — Specifies the TTL for a returned linktrace. Values 0 to 255

loopback

Syntax	loopback { <i>mac-address</i> multicast } mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [send-count <i>send-count</i>] [size <i>data-size</i>] [lbm-padding <i>padding-size</i>] [priority <i>priority</i>] [interval <i>interval-time</i>] [timeout <i>timeout</i>]
Context	oam>eth-cfm
Description	The command specifies to initiate a loopback test.
Parameters	<i>mac-address</i> — Specifies a unicast MAC address or multicast MAC address. The last nibble of the mcast address must match the level of the local MEP, or the command will error and the test will not be instantiated. multicast — Builds the class one destination multicast address based on the level of the local MEP. The last nibble of the multicast address must match the level of the local MEP or the command will error and the test will not be instantiated. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295

association *ma-index* — Specifies the MA index.

Values 1 to 4294967295

send-count *send-count* — Specifies the number of messages to send, expressed as a decimal integer. Loopback messages are sent back to back, with no delay between the transmissions.

Values 1 to 1024

Default 1

size *data-size* — This is the size of the data portion of the data TLV allowing for an optional octet string to be specified. If 0 is specified no data TLV is added to the packet. This is mutually exclusive with **lbm-padding**.

Values 0 to 1500

Default 0

lbm-padding *padding-size* — This is the size of the data portion of the data TLV and does not allow for an optional octet string. MSDU will not be processed when lbm-padding is in use. If 0 is specified, no data TLV is added to the packet. This is specified with an octet string. This is mutually exclusive with **size**.

Values 0 | 3 to 9000

Default 0

priority *priority* — Specifies a 3-bit value to be used in the VLAN tag, if present, in the transmitted frame.

Values 0 to 7

Default ccm-ltm-priority for the MEP (7)

interval *interval-time* — The interval parameter in deciseconds (100 ms) increments. This parameter is used to configure the spacing between probes within the test run. A value of 0 means probes will be sent with no enforced delay. This value is only applicable to tests where the **send-count** is 5 or less.

Values 0 to 600

Default 0 or 10 depending on send-count

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 to 10

Default 5

eth-test

Syntax	eth-test <i>mac-address mep mep-id domain md-index association ma-index [priority priority]</i> <i>[data-length data-length]</i>
Context	oam>eth-cfm
Description	This command issues an ETH-CFM test.
Parameters	<i>mac-address</i> — Specifies a unicast MAC address. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 data-length <i>data-length</i> — Indicates the UDP data length of the echo reply, the length starting after the IP header of the echo reply. Values 64 to 1500 Default 64

one-way-delay-test

Syntax	one-way-delay-test <i>mac-address mep mep-id domain md-index association ma-index [priority priority]</i>
Context	oam>eth-cfm
Description	This command issues an ETH-CFM one-way delay test.
Parameters	<i>mac-address</i> — Specifies a unicast MAC address. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 priority <i>priority</i> — Specifies the priority. Values 0 to 7 Default The CCM and LTM priority of the MEP.

two-way-delay-test

Syntax	two-way-delay-test <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [priority <i>priority</i>]
Context	oam>eth-cfm
Description	This command issues an ETH-CFM two-way delay test.
Parameters	<i>mac-address</i> — Specifies a unicast MAC address. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 priority <i>priority</i> — Specifies the priority. Values 0 to 7 Default The CCM and LTM priority of the MEP.

two-way-slm-test

Syntax	two-way-slm-test <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [priority <i>priority</i>] [send-count <i>send-count</i>] [size <i>data-size</i>] [timeout <i>timeout</i>] [interval <i>interval</i>]
Context	oam>eth-cfm
Description	This command configures an Ethernet CFM two-way SLM test in SAA.
Parameters	<i>mac-address</i> — Specifies a unicast destination MAC address. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 priority <i>priority</i> — Specifies the priority. Values 0 to 7

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. The message interval value must be expired before the next message request is sent.

Values 1 to 1000

Default 1

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 to 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 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 to 10

interval *interval* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to configure the spacing between probes within a test run.

Default 5

Values 0.1 to 0.9 | 1 to 10

alarm-notification

Syntax **alarm-notification fng-alarm-time [time] fng-reset-time [time]**

Context

```

config>service>vpls>eth
config>service>epipe>sap>eth-cfm>mep
config>service>epipe>sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>ies>if>sap>eth-cfm>mep
config>service>ies>if>spoke-sdp>eth-cfm>mep
config>service>ies>sub-if>grp-if>sap>eth-cfm>mep
config>service>vprn>if>sap>eth-cfm>mep
config>service>vprn>if>spoke-sdp>eth-cfm>mep
config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep

```

```
config>service>ipipe>sap>eth-cfm>mep
config>port>ethernet>eth-cfm>mep
config>lag>eth-cfm>eth-cfm>mep
config>router>if>eth-cfm>mep
```

Description	This command allows the operator to configure the Fault Notification Generation time values for raising the alarm and resetting the ccm defect alarm. These timers are used for network management processes and are not tied into delaying the notification to the fault management system on the network element. These timers will not affect fault propagation mechanisms.
Parameters	fng-alarm-time <i>time</i> — Specifies the time in centiseconds (10ms intervals) a defect condition at or above the low-priority-defect must be present before raising alarm. Values 0,250,500,1000 Default 0 fng-rest-time <i>time</i> — Specifies the time in centiseconds (10ms intervals) a defect condition at or above the low-priority-defect must be cleared before resetting the alarm. Values 0,250,500,1000 Default 0

sender-id

Syntax	sender-id [system local <i>local-name</i>] no sender-id
Context	config>eth-cfm>system
Description	This command allows the operator to include the configured “system name” (chassis3) or a locally configured value in ETH-CFM PDUs sent from MEPs and MIPs. The operator may only choose one of these options to use for ETH-CFM. MEPs will include the sender-id TLV for CCM (not sub second CCM enabled MEPs), LBM/LBR, and LTM/LTR. MIPs will include this value in the LBR and LTR PDUs.



Note: LBR functions reflect back all TLVs received in the LBM unchanged, including the SenderID TLV.

Parameters	system — Keyword allowing ETH-CFM to use configured “system name” value as the chassis(3). local — Provides the option to configure a local string that is different from the “system name” chassis(3) value that may be used for other means.

local-name — Specifies alpha number string up to 45 characters.

Default [no] sender-id

id-permission

Syntax **id-permission** *chassis*
no id-permission

Context config>eth-cfm>domain>assoc>bridge

Description This command allows the operator to include the sender-id TLV information that was specified under the **config>eth>system> sender-id** configuration for Service MEPs and MIPs. When this option is present under the maintenance association, the specific MPs in the association will include the sender-id tlv information in ETH-CFM PDUs. MEPs will include the sender-id TLV for CCM (not sub second CCM enabled MEPs), LBM/LBR, and LTM/LTR. MIPs will include this value in the LBR and LTR PDUs.



Note: LBR functions reflect back all TLVs received in the LBM unchanged including the SenderID TLV. Transmission of the Management Domain and Management Address fields are not supported in this TLV.

Parameters *chassis* — Sends the configured chassis information defined under>eth-cfm>system using the sender-id option.

Default [no] sender-id

facility-id-permission

Syntax **facility-id-permission** *chassis*
no facility-id-permission

Context config>eth-cfm>domain>assoc

Description This command allows the operator to include the sender-id TLV information that was specified under the **config>eth>system> sender-id** configuration for facility base MEPs. When this option is present under the maintenance association, the specific MPs in the association will include the sender-id tlv information in ETH-CFM PDUs. MEPs will include the sender-id TLV for CCM (not sub second CCM enabled MEPs), LBM/LBR, and LTM/LTR. MIPs will include this value in the LBR and LTR PDUs.



Note: LBR functions reflect back all TLVs received in the LBM unchanged including the SenderID TLV. This command will produce an error when a bridge-identifier is configured under the association. Facility MEPs do not support the bridge-identifier. Transmission of the Management Domain and Management Address fields are not supported in this TLV.

Parameters *chassis* — Sends the configured chassis information defined under >eth-cfm>system using the sender-id option.

Default [no] facility-id-permission

interface-support-enable

Syntax [no] interface-support-enable

Context config>service>epipe>sap>eth-cfm>mep>ais
config>service>epipe>spoke-sdp>eth-cfm>mep>ais
config>service>vpls>sap>eth-cfm>mep>ais
config>service>vpls>spoke-sdp>eth-cfm>mep>ais
config>service>vpls>mesh-sdp>eth-cfm>mep>ais

Description This command enable the AIS function to consider the operational state of the entity on which it is configured. With this command, ETH-AIS on DOWN MEPs will be triggered and cleared based on the operational status of the entity on which it is configured. If CCM is also enabled then transmission of the AIS PDU will be based on either the non operational state of the entity or on ANY CCM defect condition. AIS generation will cease if BOTH operational state is UP and CCM has no defect conditions. If the MEP is not CCM enabled then the operational state of the entity is the only consideration assuming this command is present for the MEP.

Default [no] interface-support-enabled: AIS will not be generated or stopped based on the state of the entity on which the DOWN MEP is configured.

csf-enable

Syntax csf-enable [*multiplier multiplier-value*]
no csf-enable

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep
config>service>ies>if>sap>eth-cfm>mep
config>service>ies>if>spoke-sdp>eth-cfm>mep
config>service>ies>sub-if>grp-if>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep
config>service>vprn>if>sap>eth-cfm>mep
config>service>vprn>if>spoke-sdp>eth-cfm>mep
config>service>vprn>sub-if>grp-if>sap>eth-cfm>mep

Description This command enable Enabled the reception and local processing of ETH-CSF frames.

Parameters	multiplier <i>multiplier-value</i> — The multiplication factor applied to the receive time used to clear the CSF condition in increments of.5.
	Values 0.0,2.0, to.30.0 — Value 0 means only clear when C-DCI is received.
	Default 3.5

3.12.2.1.9 Service Assurance Agent (SAA) Commands

saa

Syntax	saa
Context	config
Description	This command creates the context to configure the Service Assurance Agent (SAA) tests.

test

Syntax	test <i>name</i> [owner <i>test-owner</i>] no test <i>name</i>
Context	config>saa
Description	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	<i>name</i> — Identifies the saa test name to be created or edited. owner <i>test-owner</i> — Specifies the owner of an SAA operation up to 32 characters in length. Values If a <i>test-owner</i> value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy

Context	config>saa>test
Description	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	<i>acct-policy-id</i> — Specifies the accounting <i>policy-id</i> as configured in the config>log>accounting-policy context. Values 1 to 99

description

Syntax	description <i>description-string</i> no description
Context	config>saa>test
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file. The no form of this command removes the string from the configuration.
Default	No description associated with the configuration context
Parameters	<i>string</i> — 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	This command specifies whether the SAA test is continuous. Once you have configured a test as continuous, you cannot start or stop it by using the oam saa test-name {start stop} command. This option is not applicable to all SAA test types. The no form of the command disables the continuous execution of the test.

jitter-event

Syntax	jitter-event rising-threshold <i>threshold</i> [falling-threshold <i>threshold</i>] [direction] no jitter-event
Context	config>saa>test
Description	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 <i>threshold</i> — 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 to 2147483 milliseconds falling-threshold <i>threshold</i> — 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 to 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 <i>threshold</i> [falling-threshold <i>threshold</i>] [direction] no latency-event
Context	config>saa>test
Description	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 <i>threshold</i> — 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 <code>tmnxOamSaaThreshold</code>, logger application OAM, event #2101. Default 0 Values 0 to 2147483 milliseconds falling-threshold <i>threshold</i> — 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 <code>tmnxOamSaaThreshold</code>, logger application OAM, event #2101. Default 0 Values 0 to 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 <i>threshold</i> [falling-threshold <i>threshold</i>] [direction] no loss-event
Context	config>saa>test
Description	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 <i>threshold</i> — 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 to 2147483647 packets falling-threshold <i>threshold</i> — 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 to 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

trap-gen

Syntax trap-gen

Context	config>saa>test
Description	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	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 0..15
Context	config>saa>test>trap-gen
Description	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

probe-history

Syntax	probe-history [auto drop keep]
Context	config>saa>test
Description	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.

test-completion-enable

Syntax	[no] test-completion-enable
Context	config>saa>test>trap-gen
Description	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	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 0..15
Context	config>saa>test>trap-gen
Description	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	[no] type
Context	config>saa>test
Description	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

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>] [profile [in out]] [interval <i>interval</i>] [send-count <i>send-count</i>] [send-control]				
Context	oam config>saa>test>type				
Description	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 <table> <tr> <td><i>service-id:</i></td><td>1 to 2147483647</td></tr> <tr> <td><i>svc-name:</i></td><td>64 characters maximum</td></tr> </table> 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> — Specifies an unused IP address in the same network that is associated with the VPLS or PBB Epipe. ttl <i>vc-label-ttl</i> — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer. Values 1 to 255 Default 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. This parameter is only valid for VPLS services. Default MAC OAM reply sent using the data plane.	<i>service-id:</i>	1 to 2147483647	<i>svc-name:</i>	64 characters maximum
<i>service-id:</i>	1 to 2147483647				
<i>svc-name:</i>	64 characters maximum				

source-mac *ieee-address* — Specifies 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 or CFM is used. This parameter is not applicable to CPE ping on Epipe.

fc *fc-name* — The forwarding class of the MPLS echo request encapsulation.

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

Default be

profile [*in* | *out*] — The profile state of the MPLS echo request encapsulation for VPLS and the ARP packet for PBB Epipe and Epipe VLLs.

Default out

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.

Values 1 to 10

Default 1

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.

Values 1 to 100

Default 1

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane. This parameter is only valid for VPLS services.

Default MAC OAM request sent using the data plane.

dns

Syntax	dns <i>target-addr</i> <i>dns-name</i> name-server <i>ip-address</i> [source <i>ip-address</i>] [send-count <i>send-count</i>] [time-out <i>timeout</i>] [interval <i>interval</i>]
Context	oam config>saa>test>type
Description	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.

source *ip-address* — Specifies the IP address to be used as the source for performing an OAM ping operation.

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 to 100

time-out *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 to 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 to 10

eth-cfm-linktrace

Syntax	eth-cfm-linktrace <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [ttl <i>tll-value</i>] [fc { <i>fc-name</i> }] [profile { in out }] [send-count <i>send-count</i>] [timeout <i>interval</i>] [interval <i>interval</i>]
Context	config>saa>test>type
Description	This command configures a CFM linktrace test in SAA.
Parameters	<i>mac-address</i> — Specifies a unicast destination MAC address.

mep *mep-id* — Specifies the local mep-id.

Values 1 to 8191

domain *md-index* — Specifies the MD index.

Values 1 to 4294967295

association *ma-index* — Specifies the MA index.

Values 1 to 4294967295

ttl *ttl-value* — Specifies the maximum number of hops traversed in the linktrace.

Values 1 to 255

Default 64

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.

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

Default nc

profile {**in** | **out**} — The profile state of the MPLS echo request encapsulation.

Default in

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.

Values 1 to 10

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 to 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.

Values 1 to 10

Default 5

eth-cfm-loopback

Syntax	eth-cfm-loopback <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [size <i>data-size</i>] [fc { <i>fc-name</i> }] [profile { <i>in out</i> }] [send-count <i>send-count</i>] [time-out <i>interval</i>] [interval <i>interval</i>]
Context	config>saa>test>type
Description	This command configures an Ethernet CFM loopback test in SAA.
Parameters	<i>mac-address</i> — Specifies a unicast destination MAC address. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 size <i>data-size</i> — 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 to 1500 Default 0 fc <i>fc-name</i> — 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. Values be, l2, af, l1, h2, ef, h1, nc Default nc profile {<i>in out</i>} — The profile state of the MPLS echo request encapsulation. Default in 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 to 100

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 to 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.

Values 1 to 10

Default 5

eth-cfm-two-way-delay

Syntax	eth-cfm-two-way-delay <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [fc { <i>fc-name</i> }] [profile { <i>in out</i> }] [send-count <i>send-count</i>] [time-out <i>interval</i>] [interval <i>interval</i>]
Context	config>saa>test>type
Description	This command configures an Ethernet CFM two-way delay test in SAA.
Parameters	<i>mac-address</i> — Specifies a unicast destination MAC address. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 fc <i>fc-name</i> — 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. Values be, l2, af, l1, h2, ef, h1, nc Default 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. The message interval value must be expired before the next message request is sent.

Values 1 to 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 to 10

Default 5

interval *interval* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to configure the spacing between probes within a test run.

Values 1 to 10

eth-cfm-two-way-slm

Syntax	eth-cfm-two-way-delay <i>mac-address</i> mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [fc <i>{fc-name}</i>] [send-count <i>send-count</i>] [size <i>data-size</i>] [timeout <i>timeout</i>] [interval <i>interval</i>]
Context	config>saa>test>type
Description	This command configures an Ethernet CFM two-way SLM test in SAA.
Parameters	<i>mac-address</i> — Specifies a unicast destination MAC address. mep <i>mep-id</i> — Specifies the local mep-id. Values 1 to 8191 domain <i>md-index</i> — Specifies the MD index. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 fc <i>fc-name</i> — 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. Default nc Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out} — The profile state of the MPLS echo request encapsulation.

Default in

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. The message interval value must be expired before the next message request is sent.

Default 1

Values 1 to 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 to 1500

timeout *timeout* — 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.

Default 5

Values 1 to 10

interval *interval* — The interval parameter in seconds, expressed as a decimal integer. This parameter is used to configure the spacing between probes within a test run.

Values 1 to 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*} | {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance* | **service-name** *service-name*] [**time-out** *interval*]

Context config>saa>test>type

Description 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
	x:x:x:x:x:x.d.d.d.d
x:	[0 to FFFF]H

d: [0 to 255]D



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

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.

Values

ipv6-address: x:x:x:x:x:x:x[-interface]
 x:x:x:x:x:x.d.d.d.d[-interface]
 x: [0 to FFFF]H
 d: [0 to 255]D
 interface (32 chars max, mandatory for link local addresses)



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

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 to 128

tos *type-of-service* — Specifies the service type.

Values 0 to 255

size *bytes* — The request packet size in bytes, expressed as a decimal integer.

Values 0 to 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 to 65535

source *ip-address* | *dns-name* — Specifies the IP address to be used.

Values

ipv4-address: a.b.c.d
 ipv6-address: x:x:x:x:x:x:x
 x:x:x:x:x:x.d.d.d.d

	x:	[0 to FFFF]H
	d:	[0 to 255]D
dns-name:	128 characters	
	max	



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

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.

Values 1 to 10

Default 1

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)
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
x:	[0 to FFFF]H
d:	[0 to 255]D



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

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 to 100000

Default 5

do-not-fragment — Sets the DF (Do Not Fragment) bit in the ICMP ping packet (does not apply to ICMPv6).

router *router-instance* — Specifies the router name or service ID.

Values

router-name: Base, management
service-id: 1 to 2147483647

Default Base

service-name *service-name* — Specifies the service name as an integer or string.

Values

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

timeout *timeout* — 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 to 10

Default 5

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 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
x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

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 to 255

wait *milliseconds* — The time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Values 1 to 60000

tos *type-of-service* — Specifies the service type.

Values 0 to 255

Default 5000

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
 x:x:x:x:x:d.d.d.d
 x: [0 to FFFF]H
 d: [0 to 255]D



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

router *router-instance* — Specifies the router name or service ID.

Values

router-name: Base, management
service-id: 1 to 2147483647

Default Base

Isp-ping

Syntax **Isp-ping** *Isp-name* [**path** *path-name*]
Isp-ping **bgp-label** **prefix** *ip-prefix/mask* [**path-destination** *ip-address* [**interface** *if-name* | **next-hop** *ip-address*]]
Isp-ping **prefix** *ip-prefix/mask* [**path-destination** *ip-address* [**interface** *if-name* | **next-hop** *ip-address*]]
Isp-ping **static** *Isp-name* [**assoc-channel** *ipv4* | *non-ip* | *none*][**dest-global-id** *global-id* **dest-node-id** *node-id*] [**path-type** *active* | *working* | *protect*]

NOTE: Options common to all lsp-ping cases: [**fc** *fc-name* [**profile** *-[in | out]*] [**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	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 64 characters long. dest-global-id <i>global-id</i> — 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. Values 0 to 4294967295 Default 0 dest-node-id <i>node-id</i> — 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. Values a.b.c.d 1 to 4294967295> Default 0

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
 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 the MPLS-TP PSC protocol for sending user plane traffic. If MPLS-TP linear protection is not configured, then this will be the working path.

working — The working path of the MPLS-TP LSP.

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-address — 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
ipv6-address:	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
x:	[0 to FFFF]H
d:	[0 to 255]D



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

fc *fc-name* — The fc and profile parameters are used to indicate the forwarding class and profile 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 and profile 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 and profile 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 16](#) summarizes this behavior.

Table 16 **Isp-ping Request Packet and Behavior**

cpm (sender node)	echo request packet: • packet{tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
outgoing interface (sender node)	echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	echo request packet: • packet{tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
cpm (responder node)	echo reply packet: • packet{tos=1, fc2, profile2}

Table 16 **Isp-ping Request Packet and Behavior (Continued)**

outgoing interface (responder node)	echo reply packet: • pkt queued as {fc2, profile2} • 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)	echo reply packet: • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} 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.

Default be

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

profile [*in* | *out*] — The profile state of the MPLS echo request packet.

Default out

size *octets* — The MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeros to the specified 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

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 to 100

Default 1

time-out *timeout* — 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 to 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. When the LDP FEC prefix is IPv6, the user must enter a 127/8 IPv4 mapped IPv6 address, that is, in the range ::ffff:127/104.

interface *if-name* — Specifies the name of an IP interface to send the MPLS echo request message to. The name must already exist in the **config>router>interface** context.

next-hop *ip-address* — Specifies the next-hop address to send the MPLS echo request message to.

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:x.d.d.d.d
 x: [0 to FFFF]H
 d: [0 to 255]D

prefix *ip-prefix/mask* — Specifies the address prefix and subnet mask of the target LDP FEC.

Values

<ipv4-prefix>/32 | <ipv6-prefix>/128
 ipv4-prefix - a.b.c.d
 ipv6-prefix - x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:x.d.d.d.d
 x - [0 to FFFF]H
 d - [0 to 255]D

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 *none* | *non-ip* — Specifies the launched echo request's usage of the Associated Channel (ACH) mechanism, when testing an MPLS-TP LSP.

Values *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.

Output

Sample Output

This sample output is for a LDP IPv4 and IPv6 prefix FECs.

```
A:Dut-C# 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:Dut-C#

*A:Dut-A# oam lsp-ping prefix fc00::a14:106/128

LSP-PING fc00::a14:106/128: 116 bytes MPLS payload

Seq=1, send from intf A_to_B, reply from fc00::a14:106

udp-data-len=32 ttl=255 rtt=7.16ms rc=3 (EgressRtr)

---- LSP fc00::a14:106/128 PING Statistics ----
```

```

1 packets sent, 1 packets received, 0.00% packet loss

round-trip min = 7.16ms, avg = 7.16ms, max = 7.16ms, stddev = 0.000ms

*A:Dut-A#

```

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* | *non-ip* | *none*][**dest-global-id** *global-id* **dest-node-id** *node-id*] [**path-type** *active* | *working* | *protect*]

NOTE: Options common to all **lsp-trace** cases: [**detail**] [**downstream-map-tlv** {**dsmap** | **ddmap** | **none**}] [**fc** *fc-name* [**profile** *in* | *out*]] [**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 The **lsp-trace** command performs an LSP traceroute using the protocol and data structures defined in IETF RFC 4379.

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 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.

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.

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.
- 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
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
- prefix ip-prefix/mask** — Specifies the address prefix and subnet mask of the target LDP FEC.
- Values**
- | | |
|--------------------------------------|---------------------------------------|
| <ipv4-prefix>/32 <ipv6-prefix>/128 | |
| ipv4-prefix | - a.b.c.d |
| ipv6-prefix | - x:x:x:x:x:x:x (eight 16-bit pieces) |
| | x:x:x:x:x:d.d.d.d |
| x - | [0 to FFFF]H |
| d - | [0 to 255]D |

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

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 ipv4-address: a.b.c.d

ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:x.d.d.d.d

x - [0 to FFFF]H

d - [0 to 255]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 to 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 to 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 to 255

Default 5

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 to 100

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 to 10

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 to 10

Default 1

fc *fc-name* — The **fc** and **profile** parameters are used to indicate the forwarding class and profile 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 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** and **profile** 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** and **profile** 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 17](#) summarizes this behavior.

Table 17 **Isp-trace Request Packet and Behavior**

cpm (sender node)	echo request packet: • packet{tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
outgoing interface (sender node)	echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	echo request packet: • packet{tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
cpm (responder node)	echo reply packet: • packet{tos=1, fc2, profile2}
outgoing interface (responder node)	echo reply packet: • pkt queued as {fc2, profile2} • 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)	echo reply packet: • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface

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

Default be

profile in | out — The profile state of the MPLS echo request packet.

Default out

path-destination ip-address — Specifies the IP address of the path destination from the range 127/8. When the LDP FEC prefix is IPv6, the user must enter a 127/8 IPv4 mapped IPv6 address, that is, in the range ::ffff:127/104.

interface *interface-name* — Specifies the name of an IP interface to send the MPLS echo request to. The name must already exist in the con-fig>router>interface context.

next-hop *ip-address* — Specifies the next-hop to send the MPLS echo request message to.

Values

ipv4-address: a.b.c.d (host bits must be 0)

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

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 | ddmap}**.

Output

Sample Output

```
*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv ddmap 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:Dut-C#
A:Dut-C# 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:Dut-C#

```

lsp-trace over an unnumbered IP interface

```

*A:Dut-A# oam lsp-trace prefix 10.20.1.6/32 downstream-map-tlv ddmapp 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

lsp-trace of a LDP IPv6 prefix FEC

```

```

*A:Dut-A# oam lsp-trace prefix fc00::a14:106/128 path-destination ::ffff:127.0.0.1

lsp-trace to fc00::a14:106/128: 0 hops min, 0 hops max, 224 byte packets

1 fc00::a14:102 rtt=1.61ms rc=8 (DSRtrMatchLabel) rsc=1

2 fc00::a14:103 rtt=3.51ms rc=8 (DSRtrMatchLabel) rsc=1

3 fc00::a14:104 rtt=4.65ms rc=8 (DSRtrMatchLabel) rsc=1

```

```

4  fc00::a14:106  rtt=7.02ms rc=3(EgressRtr) rsc=1

*A:Dut-A# oam lsp-trace prefix fc00::a14:106/128 path-destination ::ffff:127.0.0.2
lsp-trace to fc00::a14:106/128: 0 hops min, 0 hops max, 224 byte packets

1  fc00::a14:102  rtt=1.90ms rc=8(DSRtrMatchLabel) rsc=1
2  fc00::a14:103  rtt=3.10ms rc=8(DSRtrMatchLabel) rsc=1
3  fc00::a14:105  rtt=4.61ms rc=8(DSRtrMatchLabel) rsc=1
4  fc00::a14:106  rtt=6.45ms rc=3(EgressRtr) rsc=1

```

mac-ping

Syntax	mac-ping service <i>service-id</i> destination <i>dst-ieee-address</i> [source <i>src-ieee-address</i>] [fc <i>fc-name</i> [profile in out]] [size <i>octets</i>] [ttl <i>vc-label-ttl</i>] [send-count <i>send-count</i>] [send-control] [return-control] [interval <i>interval</i>] [time-out <i>interval</i>]
Context	oam config>saa>test>type
Description	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 to 2147483647
svc-name: 64 characters maximum

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.

Values 1 to 65535
Default No OAM packet padding

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Values 1 to 255
Default 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 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.

Values Any unicast MAC value
Default The system MAC address

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.

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.

Values 1 to 100

Default 1

timeout *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 to 10

Default 5

sdp-ping

Syntax	sdp-ping orig-sdp-id [resp-sdp resp-sdp-id] [fc fc-name [profile {in out}]] [time-out interval] [interval interval] [size octets] [send-count send-count]
Context	oam config>saa>test>type
Description	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 <i>service-id</i> 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 [Table 18](#) shows the response messages sorted by precedence.

Table 18 **spd-ping Response Messages**

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

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.

Table 19 **spd-ping Test Results**

Field	Description	Values
Request Result	The result of the spd-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 .	orig-sdp-id
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
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.	orig-path-mtu
		N/A
Responding SDP-ID	The SDP-ID requested as the far-end path to respond to the spd-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.	resp-sdp-id
		N/A
Responding SDP-ID Path Used	Displays whether the responding router used the responding sdp-id to respond to the spd-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 sdp-id as the return path, Yes will be displayed. If the far-end does not use the responding sdp-id as the return path, No will be displayed. If resp-sdp is not specified, N/A will be displayed.	Yes
		No
		N/A

Table 19 **spd-ping Test Results (Continued)**

Field	Description	Values
Responding SDP-ID Administrative State	The administrative state of the responding sdp-id . 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 router 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
Responding SDP-ID Operational State	The operational state of the far-end sdp-id associated with the return path for <i>service-id</i> . When a return path is operationally down, Oper-Down is displayed. If the return sdp-id is operationally up, Oper-Up is displayed. If the responding sdp-id 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	resp-path-mtu
		N/A
Local Service IP Address	The local system IP address used to terminate remotely configured sdp-ids (as the sdp-id far-end address). If an IP address has not been configured to be the system IP address, N/A is displayed.	system-ip-addr
		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.	system-interface-name
		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> .	orig-sdp-far-end-addr
		dest-ip-addr
		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.	resp-ip-addr
		N/A
Responders Expected Far End Address	The expected source of the originators sdp-id from the perspective of the remote router terminating the sdp-id . If the far-end cannot detect the expected source of the ingress sdp-id , N/A is displayed.	resp-rec-tunnel-far-end-addr
		N/A

Table 19 **spd-ping Test Results (Continued)**

Field	Description	Values
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.	delta-request-reply
		N/A

Parameters *orig-sdp-id* — The SDP-ID to be used by **spd-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/GRE 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, **spd-ping** will attempt to send the next request if required).

Values 1 to 17407

resp-sdp *resp-sdp-id* — Optional parameter is used to specify the return SDP-ID to be used by the far-end router for the message reply for round trip SDP connectivity testing. If *resp-sdp-id* does not exist on the far-end router, terminates on another router different than the originating router, or another issue prevents the far-end router from using *resp-sdp-id*, the SDP echo reply will be sent using generic IP/GRE 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.

Values 1 to 17407

Default null. Use the non-SDP return path for message reply.

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 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 DSCP or LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating router. This is displayed in the response message output upon receipt of the message reply.

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

Default be

profile {in | out} — The profile state of the SDP encapsulation.

Default out

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 to 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.

Values 1 to 10

Default 1

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/GRE 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.

Values 40 to 9198

Default 40

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.

Values 1 to 100

Default 1

Output

Single Response Round Trip Connectivity Test Sample Output

```
A:router1> sdp-ping 10 resp-sdp 22 fc ef
Request Result: Sent - Reply Received
RTT:30ms
```

Err	SDP-ID Info	Local	Remote
—	SDP-ID:	10	22
—	Administrative State:	Up	Up
—	Operative State:	Up	Up

```

— Path MTU 4470 4470
— Response SDP Used: Yes

Err System IP Interface Info
Local Interface Name: "ESR-System-IP-Interface (Up to 32 chars)..."
— Local IP Interface State: Up
— Local IP Address: 10.10.10.11
— IP Address Expected By Remote: 10.10.10.11
— Expected Remote IP Address: 10.10.10.10
— Actual Remote IP Address: 10.10.10.10

Err FC Mapping Info Local Remote
— Forwarding Class Assured Assured
— Profile In In

```

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.

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:router1> sdp-ping 6 resp-sdp 101size 1514 count 5
Request      Response      RTT
-----
      1      Success      10ms
      2      Success      15ms
      3      Success      10ms
      4      Success      20ms
      5      Success      5ms
Sent: 5      Received: 5
Min: 5ms      Max: 20ms      Avg: 12ms

```

vccv-ping

Syntax **vccv-ping** *sdp-id:vc-id* [**target-fec-type** **static-pw-fec** *agi agi-value* **pw-path-id-saii** *src-global-id:src-node-id:src-ac-id* **pw-path-id-taii** *dest-global-id:dest-node-id:dest-ac-id*] [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*] [**reply-mode** {**ip-routed**|**control-channel**}] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**ttl** *vc-label-ttl*]

vccv-ping spoke-sdp-fec *spoke-sdp-fec-id* [**saii-type2** *global-id:prefix:ac-id* **taii-type2** *global-id:prefix:ac-id*] [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr*] [**reply-mode** {**ip-routed**|**control-channel**}] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*]

```
[timeout timeout] [interval interval] [ttl vc-label-ttl]
vccv-ping saii-type2 global-id:prefix:ac-id taii-type2 global-id:prefix:ac-id [src-ip-address
ip-addr dst-ip-address ip-addr] [reply-mode {ip-routed|control-channel}] [fc fc-name
[profile {in|out}]] [size octets] [count send-count] [timeout timeout] [interval interval] [ttl vc-
label-ttl]
vccv-ping static sdp-id:vc-id [target-fec-type pw-id-fec sender-src-address ip-address
remote-dst-address ip-address pw-id value pw-type value ] [dest-global-id global-id dest-
node-id node-id] [assoc-channel ipv4 | non-ip] [src-ip-address ip-addr] [count send-
count] [fc fc-name] [profile in|out]] [interval interval] [size octets] [timeout timeout] [ttl vc-
label-ttl][detail]
```

Context oam
config>saa>test

Description 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 both an MPLS and a GRE SDP.

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 spoke-sdp-fec id parameter. However, if the ping is across two or more segments, at least the spoke-sdp-fec id, src-ip-address ip-addr, dst-ip-address ip-addr, ttl vc-label-ttl parameters are used where:

- The src-ip-address is system IP address of the router preceding the destination router.
- 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.

Note that if the saii-type-2 and taii-type-2 parameters are specified by the user of this command for a FEC129 pseudowire, then these values will be used by the vccv-ping echo request message instead of the saii and taii of the spoke-sdp indexed by the spoke-sdp-fec parameter, or any saii and taii received in a switching point TLV for the pseudowire. Furthermore, the user must enter the saii and taii in accordance with the direction of the

pseudowire as seen from the node on which the vccv-ping command is executed. However, the values of the saii and taii sent in the echo request message will be swapped with respect to the user-entered values to match the order in the installed FEC on the targeted node. The output of the command for FEC129 type 2 pseudowire will reflect the order of the saii and taii stored on the targeted node.

This command, when used with the static option, configures a Virtual Circuit Connectivity Verification (VCCV) ping test for static MPLS-TP pseudowires used in a VLL service. It checks to verify that the destination (target) PE is the egress for the Static PW FEC. It provides for a cross-check between the dataplane and the configuration. The **vccv-ping static** command reuses an lsp-ping message format and can be used to test an MPLS-TP pseudowire VLL configured over an MPLS SDP. VCCV Ping for MPLS-TP pseudowires always uses the VCCV control word (associated channel header) with either an IPv4 channel type (0x0021) or on-demand CV message channel type (0x0025).

Note that vccv-ping static can only be initiated on a T-PE. Both the echo request and reply messages are sent using the same, in-band, encapsulation. If the target-fec-type option is not specified, then the target FEC stack contains a static PW FEC TLV. The contents of this TLV are populated based on the source Node ID, source Global ID, and Destination Global ID and Destination Node ID in the **vccv-ping** command (or taken from the pseudowire context if omitted from the command).

The target-fec-type option allows the user to test a segment of a MS-PW that does not have the same FEC type as the local segment from the T-PE where the **vccv-ping** command is issued. This is applicable for performing VCCV Ping on an MS-PW comprised of static PW FEC segments and dynamically signaled PW ID FEC segments.

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 *sdp-id:vc-id* — If a FEC 128 PW is being tested, then its VC ID 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 to 17407:1 to 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 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.

spoke-sdp-fec is mutually exclusive with the sdp-id:vc-id parameter.

Values 1 to 4294967295

saii-type2 *global-id:prefix:ac-id* — If a FEC129 All Type 2 pseudowire is being tested, then the source attachment individual identifier (SAII) must be indicated. The saii-type2 parameter is mutually exclusive with sdp-id:vc-id.

global-id — The Global ID of the router T-PE.

Values 1 to 4,294,967,295

prefix — The prefix on the router T-PE that the spoke-SDP is associated with.

ac-id — An unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 to 4,294,967,295

taii-type2 *global-id:prefix:ac-id* — If a FEC129 All 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*.

global-id — The Global ID of the far end T-PE of the FEC129 pseudowire.

Values 1 to 4,294,967,295

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 to 4,294,967,295

src-ip-address *ip-addr* — Specifies the source IP address.

Values ipv4-address: a.b.c.d

dst-ip-address *ip-addr* — Specifies the destination IP address.

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 RFE 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 the 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 forwarding class used by the far-end 7750 SR 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.

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

Default be

The TOS byte is not modified. [Table 20](#) summarizes this behavior.

Table 20 vccv-ping Request Packet and Behavior

cpm (sender node)	echo request packet: • packet{tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
outgoing interface (sender node)	echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	echo request packet: • packet{tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
cpm (responder node)	echo reply packet: • packet{tos=1, fc2, profile2}
outgoing interface (responder node)	echo reply packet: • pkt queued as {fc2, profile2} • 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)	echo reply packet: • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface

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

Default be

profile {in | out} — The profile state of the MPLS echo request encapsulation.

Default out

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.

Values 1 to 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.

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 to 10

Default 1

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 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

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.

Values 1 to 1000

Default 1

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.

dest-global-id *global-id* — The MPLS-TP global ID for the far end node of the pseudowire under test. If this is not entered, then the dest-global-id is taken from the pseudowire context.

dest-node-id *node-id* — The MPLS-TP node ID of the far-end node for the pseudowire under test. If this is not entered, then the dest-global-id is taken from the pseudowire context.

assoc-channel {ipv4 | non-ip} — The associated channel encapsulation format to use for the VCCV ping echo request and echo reply packet for a PW that uses the static PW FEC. An associated channel type of ipv4 must be used if a vccv-ping is performed to a remote segment of a different FEC type.

Values **ipv4** — IPv4 encapsulation in an IPv4 pseudowire associated channel (channel type 0x0021)
 non-ip —MPLS-TP encapsulation without UDP/IP headers, in pseudowire associated channel using channel type 0x025.

Default non-ip

target-fec-type {pw-id-fec | static-pw-fec} — The FEC type for a remote PW segment targeted by a VCCV Ping echo request. This parameter is used if VCCV Ping is used along a MS-PW where a static MPLS-TP PW segment using the static PW FEC is switched to a T-LDP signaled segment using the PW ID FEC (FEC128), or vice versa, thus requiring the user to explicitly specify a target FEC that is different from the local segment FEC.

Values **pw-id-fec** — Indicates that FEC element for the remote target PW
 0000000000000000000000000000000000000000000000000000000000000000
 0000000000000000000000000000000000000000000000000000000000000000seg
 ment is of type PW ID (FEC128).
 static-pw-fec — Indicates that FEC element for the remote target PW segment is of type Static PW FEC.

agi agi-value — The attachment group identifier for the target FEC. This parameter is only valid in combination with the target-fec-type static-pw-fec.

Values 0 to 4,294,967,295

pw-path-id-saii src-global-id:src-node-id:src-ac-id — The SAI of the target FEC. This parameter is only valid in combination with the target-fec-type static-pw-fec.

src-global-id — The Global ID of the SAI of the targeted static PW FEC element.

Values 1 to 4294967295

src-node-id — The node-id on far end T-PE that the pseudowire being tested is associated with.

Values ipv4-formatted address: a.b.c.d

src-ac-id — An unsigned integer representing a locally unique SAI for the pseudowire being tested at the far end T-PE.

Values 1 to 4294967295

pw-path-id-taii dst-global-id:dst-node-id:dst-ac-id — The SAI of the target FEC. This parameter is only valid in combination with the target-fec-type static-pw-fec.

dst-global-id — The Global ID of the TAIL of the targeted static PW FEC element.

Values 1 to 4294967295

dst-node-id — The node-id of the TAIL on far end T-PE that the pseudowire being tested is associated with.

Values ipv4-formatted address: a.b.c.d

dst-ac-id — An unsigned integer representing a locally unique TAIL for the pseudowire being tested at the far end T-PE.

Values 1 to 4294967295

remote-dst-address *ipv4-address* — The 4-octet IPv4 address of the far end node that is a target of the VCCV Ping echo request. This parameter is only valid in combination with the target-fec-type static-pw-fec.

Values ipv4-formatted address: a.b.c.d

sender-src-address *ipv4-address* — The 4-octet IPv4 address of the node originating the VCCV Ping echo request. This parameter is only valid in combination with the target-fec-type pw-id.

Values ipv4-formatted address: a.b.c.d

remote-dst-address *ipv4-address* — The 4-octet IPv4 address of the far end node that is a target of the VCCV Ping echo request. This parameter is only valid in combination with the target-fec-type pw-id.

Values ipv4-formatted address: a.b.c.d

pw-type *value* — The PW Type value of the PW segment targeted on the far end node. This field must be included to populate the PW type field of the PW ID FEC in the ratert FEC static TLV, when the far end FEC type is different from the local FEC type and the target-fec-type is pw-id-fec.

Values atm-cell, atm-sdu, atm-vcc, atm-vpc, cesopsn, cesopsn-cas|ether, satop-e1, satop-t1, [1 to 65535]

Output

Sample Output

```
Ping TPE to SPE on a LDP/GRE tunnel
=====

*A:Dut-B# oam vccv-ping 3:1
VCCV-PING 3:1 88 bytes MPLS payload
Seq=1, send from intf toSPE1-D-8 to NH 12.1.8.2
      reply from 4.4.4.4 via Control Channel
      udp-data-len=56 rtt=0.689ms rc=8 (DSRtrMatchLabel)

---- VCCV PING 3:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 0.689ms, avg = 0.689ms, max = 0.689ms, stddev = 0.000ms
```

```
Ping TPE to SPE on a RSVP tunnel
=====

A:Dut-C# oam vccv-ping 5:1
VCCV-PING 5:1 88 bytes MPLS payload
Seq=1, send from intf toSPE2-E-5 to NH 12.3.5.1
      send from lsp toSPE2-E-5
      reply from 5.5.5.5 via Control Channel
      udp-data-len=56 rtt=1.50ms rc=8 (DSRtrMatchLabel)

---- VCCV PING 5:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 1.50ms, avg = 1.50ms, max = 1.50ms, stddev = 0.000ms


Ping TPE to TPE over multisegment pseudowire
=====
*A:Dut-C# oam vccv-ping 5:1 src-ip-address 4.4.4.4 dst-ip-address 2.2.2.2 pw-
id 1 ttl 3
VCCV-PING 5:1 88 bytes MPLS payload
Seq=1, send from intf toSPE2-E-5 to NH 12.3.5.1
      send from lsp toSPE2-E-5
      reply from 2.2.2.2 via Control Channel
      udp-data-len=32 rtt=2.50ms rc=3 (EgressRtr)

---- VCCV PING 5:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 2.50ms, avg = 2.50ms, max = 2.50ms, stddev = 0.000ms


Ping SPE to TPE (over LDP tunnel)
=====

Single segment:
-----

*A:Dut-D# oam vccv-ping 3:1 reply-mode ip-routed
VCCV-PING 3:1 88 bytes MPLS payload
Seq=1, send from intf toTPE1-B-8 to NH 12.1.8.1
      reply from 2.2.2.2 via IP
      udp-data-len=32 rtt=1.66ms rc=3 (EgressRtr)

---- VCCV PING 3:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 1.66ms, avg = 1.66ms, max = 1.66ms, stddev = 0.000ms


Multisegment:
-----
*A:Dut-D>config>router# oam vccv-ping 4:200 src-ip-address 5.5.5.5 dst-ip-
address 3.3.3.3 pw-id 1 ttl 2 reply-mode ip-routed
VCCV-PING 4:200 88 bytes MPLS payload
Seq=1, send from intf toSPE2-E-5 to NH 12.2.5.2
      reply from 3.3.3.3 via IP
      udp-data-len=32 rtt=3.76ms rc=3 (EgressRtr)

---- VCCV PING 4:200 Statistics ----
```

```

1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 3.76ms, avg = 3.76ms, max = 3.76ms, stddev = 0.000ms

Ping SPE to SPE
=====
*A:Dut-D# oam vccv-ping 4:200 reply-mode ip-routed
VCCV-PING 4:200 88 bytes MPLS payload
Seq=1, send from intf toSPE2-E-5 to NH 12.2.5.2
      reply from 5.5.5.5 via IP
      udp-data-len=56 rtt=1.77ms rc=8 (DSRtrMatchLabel)

---- VCCV PING 4:200 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 1.77ms, avg = 1.77ms, max = 1.77ms, stddev = 0.000ms

```

vccv-trace

Syntax	vccv-trace <i>sdp-id:vc-id</i> [reply-mode ip-routed control-channel] [target-fec-type static-pw-fec agi <i>attachment-group-identifier</i> pw-path-id-saii <i>global-id:node-id:ac-id</i> pw-path-id-taii <i>global-id:node-id:ac-id</i>] vccv-trace saii-type2 <i>global-id:prefix:ac-id</i> taii-type2 <i>global-id:prefix:ac-id</i> [reply-mode ip-routed control-channel] vccv-trace spoke-sdp-fec <i>spoke-sdp-fec-id</i> [reply-mode ip-routed control-channel] [saii-type2 <i>global-id:prefix:ac-id</i> taii-type2 <i>global-id:prefix:ac-id</i>] vccv-trace static <i>sdp-id:vc-id</i> [assoc-channel ipv4 non-ip] [src-ip-address <i>ipv4-address</i>] [target-fec-type <i>pw-id-fec</i> sender-src-address <i>ipv4-address</i> remote-dst-address <i>ipv4-address</i> pw-id <i>pw-id</i> pw-type <i>pw-type</i>] options common to all vccv-trace cases: [fc <i>fc-name</i> [profile in out]] [interval <i>interval-value</i>] [max-fail <i>no-response-count</i>] [max-ttl <i>max-vc-label-ttl</i>] [min-ttl <i>min-vc-label-ttl</i>] [probe-count <i>probe-count</i>] [size <i>octets</i>] [timeout <i>timeout-value</i>]
Context	oam config>saa>test>type
Description	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.

Note that if the `saii-type-2` and `taii-type-2` parameters are specified by the user of this command for a FEC129 pseudowire, then these values will be used by the `vccv-ping echo` request message instead of the `saii` and `taii` of the spoke-sdp indexed by the `spoke-sdp-fec` parameter, or any `saii` and `taii` received in a switching point TLV for the pseudowire. Furthermore, the user must enter the `saii` and `taii` in accordance with the direction of pseudowire as seen from the node on which the `vccv-trace` command is executed. However, the values of the `saii` and `taii` sent in the echo request message will be swapped with respect to the user-entered values to match the order in the installed FEC on the targeted node. The output of the command for a FEC129 type 2 pseudowire will reflect the order of the `saii` and `taii` stored on the targeted node.

This command, when used with the `static` option, configures a Virtual Circuit Connectivity Verification (VCCV) automated trace test for static MPLS-TP pseudowires used in a VLL service. VCCV trace for MPLS-TP pseudowires always uses the VCCV control word (associated channel header) with either an IPv4 channel type (0x0021) or on-demand CV message channel type (0x0025).

Note that `vccv-trace static` can only be initiated on a T-PE. Both the echo request and reply messages are sent using the same, in-band, encapsulation. The target FEC stack contains a static PW FEC TLV. The contents of this TLV are populated based on the source Node ID, source Global ID, and Destination Global ID and Destination Node ID taken from the pseudowire context.

The `target-fec-type` option allows the user to perform a `vccv-trace` to a segment of a MS-PW that does not have the same FEC type as the local segment from the T-PE where the `vccv-trace` command is issued. This is applicable for performing VCCV Ping on an MS-PW comprised of static PW FEC segments and dynamically signaled PW ID FEC segments.

- Parameters**
- `sdpid:vcid`** — If a FEC 128 PW is being tested, then its VC ID 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 to 17407:1 to 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 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.
`spoke-sdp-fec` is mutually exclusive with the `sdp-id:vc-id` parameter.
Values 1 to 4294967295
 - `saii-type2 global-id:prefix:ac-id`** — If a FEC129 All 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.

global-id — The global ID of this T-PE node.

Values 1 to 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 to 4294967295

taii-type2 *global-id:prefix:ac-id* — If a FEC129 All 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** parameter.

global-id — The global ID of the far end T-PE of the FEC129 pseudowire.

Values 1 to 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 to 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 SPE node. If the user attempts this, a warning is issued to use the ipv4 reply mode.

Default control-channel

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.

Values *fc-name* — The forwarding class of the VCCV trace echo request encapsulation.

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 and profile 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 and profile 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.

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

The TOS byte is not modified. [Table 21](#) summarizes this behavior.

Table 21 vccv trace Request Packet and Behavior

cpm (sender node)	echo request packet: • packet{tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
outgoing interface (sender node)	echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	echo request packet: • packet{tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
cpm (responder node)	echo reply packet: • packet{tos=1, fc2, profile2}

Table 21 vccv trace Request Packet and Behavior (Continued)

outgoing interface (responder node)	echo reply packet: • pkt queued as {fc2, profile2} • 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)	echo reply packet: • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface

profile {in | out} — The profile state of the VCCV trace echo request packet.

Default out

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 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

probe-count *probes-per-hop* — The number of VCCV trace echo request messages to send per TTL value.

Values 1 to 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 to 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 to 255

Default 1

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.

Values 1 to 255

Default 1

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.

Values 1 to 255

Default 8

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.

Values 1 to 255

Default 5

assoc-channel {*ipv4* | *non-ip*} — the associated channel encapsulation format to use for the VCCV trace echo request and echo reply packet for a PW that uses the static PW FEC. An associated channel type of *ipv4* must be used if a vccv-ping is performed to a remote segment of a different FEC type.

Values *ipv4* – IPv4 encapsulation in an IPv4 pseudowire associated channel (channel type 0x0021)

non-ip — MPLS-TP encapsulation without UDP/IP headers, in pseudowire associated channel using channel type 0x025.

Default *non-ip*

target-fec-type {*pw-id-fec* | *static-pw-fec*} — The FEC type for a remote PW segment targeted by a VCCV trace echo request. This parameter is used if VCCV trace is used along a MS-PW where a static MPLS-TP PW segment using the static PW FEC is switched to a T-LDP signaled segment using the PW ID FEC (FEC128), or vice versa, thus requiring the user to explicitly specify a target FEC that is different from the local segment FEC.

Values **pw-id-fec** — Indicates that FEC element for the remote target PW segment is of type PW ID (FEC128).

static-pw-fec — Indicates that FEC element for the remote target PW segment is of type Static PW FEC.

agi *agi-value* — The attachment group identifier for the target FEC. This parameter is only valid in combination with the target-fec-type static-pw-fec.

Values 0 to 4,294,967,295

pw-path-id-saii *src-global-id:src-node-id:src-ac-id* — The SAII of the target FEC. This parameter is only valid in combination with the target-fec-type static-pw-fec.

Values

src-global-id — The Global ID of the SAII of the targeted static PW FEC element.

Values 1 to 4,294,967,295

src-node-id — The node-id on far end T-PE that the pseudowire being tested is associated with.

Values ipv4-formatted address: a.b.c.d

src-ac-id — An unsigned integer representing a locally unique SAII for the pseudowire being tested at the far end T-PE.

Values 1 to 4,294,967,295

pw-path-id-taii *dst-global-id:dst-node-id:dst-ac-id* — The SAII of the target FEC. This parameter is only valid in combination with the target-fec-type static-pw-fec.

Values

dst-global-id — The Global ID of the TAII of the targeted static PW FEC element.

Values 1 to 4,294,967,295

dst-node-id — The node-id of the TAII on far end T-PE that the pseudowire being tested is associated with.

Values ipv4-formatted address: a.b.c.d

dst-ac-id — An unsigned integer representing a locally unique TAII for the pseudowire being tested at the far end T-PE.

Values 1 to 4,294,967,295

remote-dst-address *ipv4-address* — The 4-octet IPv4 address of the far end node that is a target of the VCCV Ping echo request. This parameter is only valid in combination with the target-fec-type static-pw-fec.

Values ipv4-formatted address: a.b.c.d

sender-src-address *ipv4-address* — The 4-octet IPv4 address of the node originating the VCCV Ping echo request. This parameter is only valid in combination with the target-fec-type pw-id.

Values ipv4-formatted address: a.b.c.d

remote-dst-address ipv4-address — The 4-octet IPv4 address of the far end node that is a target of the VCCV Ping echo request. This parameter is only valid in combination with the target-fec-type pw-id.

Values ipv4-formatted address: a.b.c.d

pw-type value — The PW Type of the PW segment targeted on the far end node. This field must be included to populate the PW type field of the PW ID FEC in the FEC static TLV, when the far end FEC type is different from the local FEC type and the target-fec-type is pw-id-fec.

Values atm-cell, atm-sdu, atm-vcc, atm-vpc, cesopsn-cas|ether, satop-e1, satop-t1, [1 to 65535].

Output

Sample Output

```
*A:138.120.214.60# oam vccv-trace 1:33
VCCV-TRACE 1:33 with 88 bytes of MPLS payload
1 1.1.63.63 rtt<10ms rc=8 (DSRtrMatchLabel)
2 1.1.62.62 rtt<10ms rc=8 (DSRtrMatchLabel)
3 1.1.61.61 rtt<10ms rc=3 (EgressRtr)
```

Trace with detail:

```
*A:138.120.214.60>oam vccv-trace 1:33 detail

VCCV-TRACE 1:33 with 88 bytes of MPLS payload
1 1.1.63.63 rtt<10ms rc=8 (DSRtrMatchLabel)
   Next segment: VcId=34 VcType=AAL5SDU Source=1.1.63.63 Remote=1.1.62.62
2 1.1.62.62 rtt<10ms rc=8 (DSRtrMatchLabel)
   Next segment: VcId=35 VcType=AAL5SDU Source=1.1.62.62 Remote=1.1.61.61
3 1.1.61.61 rtt<10ms rc=3 (EgressRtr)
SAA:

*A:multisim3>config>saa# info
-----
      test "vt1"
        shutdown
        type
          vccv-trace 1:2 fc "af" profile in timeout 2 interval 3 size 200
min-ttl 2 max-ttl 5 max-fail 2 probe-count 3
      exit
    exit
  ..
-----
*A:multisim3>config>saa#
```

3.12.2.1.10 OAM SAA Commands

saa

Syntax	saa <i>test-name</i> [owner <i>test-owner</i>] { start stop } [no-accounting]
Context	oam
Description	Use this command to start or stop an SAA test that is not configured as continuous.
Parameters	<i>test-name</i> — Name of the SAA test. The test name must already be configured in the config>saa>test context. <i>test-owner</i> — Specifies the owner of an SAA operation up to 32 characters in length. Values If a <i>test-owner</i> 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. When specifying no-accounting then 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.

3.12.2.1.11 OAM Performance Monitoring and Binning Commands

oam-pm

Syntax	oam-pm session <i>session-name</i> { dmm lmm slm twamp-light } { start stop }
Context	oam
Description	This command allows the operator to start and stop on-demand OAM-PM sessions.
Parameters	session <i>session-name</i> — Identifies the session name that the test is associated with <i>session-name</i> — Specifies the session name, up to 32 characters in length dmm — Specifies the DMM test that will be affected by the command lmm — Specifies the LMM test that will be affected by the command slm — Specifies the SLM test that will be affected by the command twamp-light — Specifies the TWAMP-light test that will be affected by the command start — Manually starts the test stop — Manually stops the test

oam-pm

Syntax	oam-pm
Context	config
Description	This is the top level context that contains the configuration parameters that defines storage parameters (including binning structures), availability/resiliency and the individual proactive, and on-demand tests used to gather the performance/statistical information.

bin-group

Syntax	bin-group <i>bin-group-number</i> [fd-bin-count <i>fd-bin-count</i> fdr-bin-count <i>fdr-bin-count</i> ifdv-bin-count <i>ifdv-bin-count</i> create]
Context	config>oam-pm
Description	This command allows the operator to configure 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 this will be assigned by default. The default values for bin-group 1 are (fd-bin-count 3 bin 1 lower-bound 5000us, bin 2 lower-bound 10000us fdr-bin-count 2 bin 1 lower-bound 5000us and ifdv-bin-count 2 bin 1 lower-bound 5000us)

Parameters	<i>bin-group-number</i> — Numerical identifier for a bin-group that is referenced by oam-pm sessions. A bin group can only shutdown and modified when all the PM Sessions referencing the bin group have been shutdown. The only exception is the description parameter. Values 1 to 255 fd-bin-count <i>fd-bin-count</i> — Specifies the number of fd bins that will be created. Values 2 to 10 fdr-bin-count <i>fdr-bin-count</i> — Specifies the number of fdr bins that will be created. Values 2 to 10 ifdv-bin-count <i>ifdv-bin-count</i> — Specifies the number of ifdv bins that will be created. Values 2 to 10 create — Keyword that instantiates the bin group.

description

Syntax	description <i>description-string</i> no description
Context	config>oam-pm>bin-group
Description	This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file. The no form of the command removes the string from the configuration
Parameters	<i>description-string</i> — 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.

bin-type

Syntax	bin-type {fd fdr ifdv}
Context	config>oam-pm>bin-group
Description	This command is the start of the hierarchy where the specific delay metric bin structure will be defined.
Parameters	fd — Keyword to enter the frame delay bin threshold configuration. fdr — Keyword to enter the frame delay range bin threshold configuration. ifdv — Keyword to enter the inter-frame delay variation bin thresholds configuration.

bin

Syntax	bin <i>bin-number</i> lower-bound <i>microseconds</i>
Context	config>oam-pm>bin-group>bin-type
Description	This command allows the operator specify the individual floors thresholds for the bins. The operator does not have to specific a lower threshold for every bin that was previously defined by the bin-count for the specific type. By default each bin will be the bin-number * 5000 microseconds. Lower thresholds in the previous adjacent bin must be lower than the threshold of the next higher bin threshold. A separate line per bin is required to configured an operator specific threshold. An error will prevent the bin from entering the active state if this is not maintained, at the time the “no shutdown” is issued. Bin 0 is the result of the difference between 0 and the configured lower-threshold 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 lower-bound removes the user configured threshold value and applies the default for the bin.
Parameters	<i>bin-number</i> — Specifies bin to configure. Values 1 to 9 lower-bound <i>microseconds</i> — The threshold that defines the floor of the bin. The bin range is the difference between its configured threshold and the threshold of the next higher bin in microsecond threshold value. Values 1 to 4294967295 Default bin-number * 5000

delay-event

Syntax	delay-event { forward backward round-trip } lowest-bin <i>bin-number</i> threshold <i>raise-threshold</i> [clear <i>clear-threshold</i>] [no] delay-event { forward backward round-trip }
Context	config>oam-pm>bin-group>bin-type
Description	This command sets the bin number, the threshold and the direction that is monitored to determine if a delay metric threshold crossing event has occurred or has cleared. It requires a bin number, a rising threshold value and a direction. If the [clear threshold] is not specified, the traffic crossing alarm will be stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. When a raise threshold is reached, the log event is generated. Each unique threshold can only be raised once for the threshold within measurement interval. If the optional clear threshold is specified, the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to

following measurement intervals. If a threshold crossing event is raised another will not be raised until a measurement interval completes, and the clear threshold has not been exceeded. A clear event will be raised under that condition. In general, alarms are generated when there is a state change. The thresholds configured will be applied to the count in specified bin and all higher number bins.

The **no** version of this command removes thresholding for this delay metric. The complete command must be configured in order to remove the specific threshold.

Default	[no] delay-events
Parameters	forward — The threshold is applied to the forward direction bin. backward — The threshold is applied to the backward direction bin. round-trip — The threshold is applied to the roundtrip direction bin. lowest-bin <i>bin-number</i> — The number of the bin that the threshold is applied to. This bin and all higher bins will be monitoring to determine if the sum total results in these bins have reached or crossed the configured threshold. Values 0 to 9 threshold <i>raise-threshold</i> — The rising numerical value in the range that determines when the event is to be generated, when value reached. Values 1 to 864000 clear <i>clear-threshold</i> — An optional numerical value in the range threshold used to indicate stateful behavior that allows the operator to configure a lower value than the rising threshold that determines when the clear event should be generated. Clear is generated when the end of measurement interval count is less than or equal to the configured value. If this option is not configured the behavior is stateless. Zero means no results can existing in the lower bin or any higher. Values 0 to 863999 Default Clear threshold disabled

delay-event-exclusion

Syntax	delay-event-exclusion { forward backward round-trip } lowest-bin <i>bin-number</i> no delay-event-exclusion { forward backward round-trip }
Context	config>oam-pm>bin-group>bin-type
Description	This optional command allows results from probes that map to the specified bin and higher bins to be excluded from the TCA count. The TCA count is used to determine if a threshold has been reached by the event monitoring function. Individual counters are incremented in the bin, but the counts in the specified bin and higher bins are not included in the TCA threshold computation. A delay-event must be configured in the same direction, and the lowest-bin configured as part of the delay-event-exclusion command must be higher than the lowest bin specified by the corresponding delay-event command.

The bin group allows this optional command to be added, modified, or deleted while tests are actively referencing the bin group. The bin group does not need to be shut down during **delay-event-exclusion** configuration. If the values are modified while the active tests are executing, all configured TCAs for the specified direction within the bin group will enter a pending (p) state until the start of the next measurement interval. Any existing stateful TCAs that were raised are cleared without creating a log event, and no further processing for the affected TCAs will occur in the active window. Depending on timing, the pending state may continue past the adjacent measurement interval until the start of the following measurement interval.

The **no** form of this command does not exclude any values from the configured TCA threshold.

Default	no delay-event-exclusion forward
	no delay-event-exclusion backward
	no delay-event-exclusion round-trip
Parameters	forward — the forward direction bin
	backward — the backward direction bin
	round-trip — the round-trip direction bin
	<i>bin-number</i> — the number of the lowest bin that the exclusion is applied to. This bin and all higher bins are excluded from the delay-event (TCA) count. If no bin numbers are configured, this command is ignored.
	Values 1 to 9

exclude-from-avg

Syntax	exclude-from-avg {forward backward round-trip} bins <i>bin-numbers</i> no exclude-from-avg {forward backward round-trip}
Context	config>oam-pm>bin-group>bin-type
Description	This optional command allows the results from probes that map to the specified bins within the bin type to be excluded from the average calculation. Individual counters are incremented in the bin, but the average is not affected by the value of the excluded delay metric for the individual probes in this bin. The bin group does not allow this command to be added, modified, or deleted when a test is actively referencing the bin group. Sessions that reference the bin group must have the bin group and tests shut down before changes can be made. The no form of this command removes the exclusion, and all bins are included in the average calculation.
Default	no exclude-from-avg forward
	no exclude-from-avg backward
	no exclude-from-avg round-trip

Parameters	forward — the forward direction bin backward — the backward direction bin round-trip — the round-trip direction bin <i>bin-numbers</i> — the bin numbers to be excluded from the average calculation. The values typically represent, but are not restricted to, the highest and lowest configured bins in order to eliminate outlying results that are not representative of network performance. A hyphen can be entered between bin numbers to include a continuous sequence of bins; for example, typing “7-9” would specify bins 7, 8, and 9. Commas can be entered between bin numbers to include separate or non-continuous bins; for example, typing “0,8,9” would specify bins 0, 8, and 9. Both hyphens and commas can be used in this manner in the same configuration; for example, typing “0,7-9” would include bins 0, 7, 8, and 9. All bin numbers specified as part of this command must be configured. If a specified bin does not exist, the command will fail. Values 0 to 9

shutdown

Syntax	[no] shutdown
Context	config>oam-pm>bin-group
Description	This command activates and deactivates the bin group. Only the description of the bin group can be modified when the bin group is in a “no shutdown” state. No other changes can be made while the bin group is active. The bin group can only be shutdown and modified when all references in the various PM Sessions or individual tests have been shutdown. If an active PM session is referencing the bin-group, it will generate an error indicating there are x number of active tests referencing the bin-group, and it cannot be shutdown. The no form of the command activates the bin group as available for PM Sessions and tests to utilize.
Default	shutdown

session

Syntax	session <i>session-name</i> test-family {ethernet ip} [session-type {proactive on-demand}] create no session <i>session-name</i>
Context	config>oam-pm

Description This command creates 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 shutdown. Individual tests sessions within the container may be shutdown. 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 shutdown. The only exception to this is the description value.

The **no** form of the command deletes the session.

Parameters *session-name* — Identifies the session container.

test-family — Indicates the type family and sets the context for the individual parameters.

ethernet — Keyword that indicates the test will be based on the Ethernet layer.

ip — Keyword that indicates the test will be based on the IP layer.

session-type — Specifies how to set the Type bit in the Flags byte, and influences how different test criteria may be applied to the individual test. Not all test-families carry this information in the PDU.

proactive — Keyword setting the type to always on with immediate start and no stop.

on-demand — Keyword setting the type a demand function with an immediate start and no stop, or stop based on offset.

Default proactive

create — Instantiates the PM session.

description

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

Context config>oam-pm>session

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

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

Parameters *description-string* — 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.

bin-group

Syntax	bin-group <i>bin-group-number</i> no bin-group
Context	config>oam-pm>session
Description	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.
Parameters	<i>bin-group-number</i> — The number that was used to create the specific bin-group that will be referenced for this session. Values 1 to 255 Default 1

meas-interval

Syntax	meas-interval {5-mins 15-mins 1-hour 1-day} create no meas-interval {5-mins 15-mins 1-hour 1-day}
Context	config>oam-pm>session
Description	This command establishes the parameters of the individual measurement intervals utilized by the session. Multiple measurement intervals may be specified within 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 15-mins 1-hour 1-day} — Keywords used to specifies the duration of the measurement interval. create — Keyword the instantiates the measurement interval.

accounting-policy

Syntax	accounting-policy <i>acct-policy-id</i> no accounting-policy
Context	config>oam-pm>session>meas-interval
Description	This optional command allows the operator to assign an accounting policy and the policy-id (configured under the config>log>accounting-policy) with a record-type of complete-pm. 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	<i>acct-policy-id</i> — 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	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. Alignment with the time of day clock always defaults to the representative top of the hour. Clock aligned 15-minute measurement intervals will divide the hour into four equal sections 00, 15, 30, 45. Clock aligned 1-hour measurement intervals will start at 00. Clock aligned 1-day measurement intervals will start at midnight. Test relative start times will launch the measurement interval when the individual test enters the active (no shutdown) state. 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 the no shutdown will exactly correspond to the clock based measurement interval start time. Clock aligned measurement intervals can include an additional offset. See clock-offset command option under this context. The no form of the command sets the boundary to the default clock-aligned.
Parameters	clock-aligned — Keyword that aligns the start of the measurement interval with the time of day clock. test-relative — Keyword that aligns the start of the measurement interval with the start of the test.

clock-offset

Syntax	clock-offset <i>seconds</i> no clock-offset
Context	config>oam-pm>session>meas-interval
Description	This command allows measurement intervals with a boundary-type of clock aligned to be offset from 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 five minutes from their default alignments with respect to the time of day clock. The no form of the command sets the offset to 0.

Parameters	<i>seconds</i> — The number of seconds to offset a clock-alignment measurement interval from its default.
Values	0 to 86399
Default	0

event-mon

Syntax	event-mon
Context	config>oam-pm>session>measurement-interval
Description	This hierarchy allows for enabling of the different threshold events on a specific measurement interval. Only one measurement interval with a configured OAM PM session can have events enabled using the no shutdown command.

delay-events

Syntax	delay-events [no] delay-events
Context	config>oam-pm>session>measurement-interval>event-monitoring
Description	This enables and disables the monitoring of all configured delay events. Adding this functionality will start the monitoring of the configured delay events at the start of the next measurement interval. If the function is removed using the no command, all monitoring of configured delay events, logging, and recording of new events for that session will be suspended. Any existing events at the time of the shutdown will be maintained until the active measurement window in which the removal was performed has completed. The state of this monitoring function can be changed without having to shutdown all the tests in the session.
Default	[no] delay-events

loss-events

Syntax	loss-events [no] loss-events
Context	config>oam-pm>session>measurement-interval>event-monitoring

Description	This enables and disables the monitoring of all configured loss events. Adding this functionality will start the monitoring of the configured loss events at the start of the next measurement interval. If the function is removed using the no command, all monitoring of configured loss events, logging, and recording of new events for that session will be suspended. Any existing events at the time of the shutdown will be maintained until the active measurement window in which the removal was performed has completed. The state of this monitoring function can be changed without having to shutdown all the tests in the session.
Default	[no] loss-events

shutdown

Syntax	[no] shutdown
Context	config>oam-pm>session>measurement-interval>event-monitoring
Description	Issuing a no shutdown command will start the monitoring of the configured events at the start of the next measurement interval. If a shutdown is issued, all monitoring of configured events, logging, and recording of new events for that session will be suspended. Any existing events at the time of the shutdown will be maintained until the active measurement window in which the event-mon shutdown was issued has completed. The state of this monitoring function can be changed without having to shutdown all the tests in the session.
Default	shutdown

intervals-stored

Syntax	intervals-stored <i>intervals</i> no intervals-stored
Context	config>oam-pm>session>meas-interval
Description	This command defines 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 memory is available. The numbers are increasing from 1 to the configured value + 1. The active pm data will be stored in the interval number 1 and older runs are stored, in order, to the upper most number with the oldest rolling off 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 get 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 both intervals are required the sum total of both cannot exceed 96. The 1-hour and 1-day measurement intervals utilizes their own ranges.

If this command is omitted when configuring the measurement interval, the default values will be used.

Parameters	<i>intervals</i> — Specifies the measurement interval.
	<i>5-mins</i> — Specifies 5 minutes measurement interval.
	Values 1 to 96
	Default 32
	<i>15-mins</i> — Specifies 15 minutes measurement interval.
	Values 1 to 96
	Default 32
	<i>1-hour</i> — Specifies 1 hour measurement interval.
	Values 1 to 24
	Default 8
	<i>1-day</i> — Specifies 1 day measurement interval.
	Values 1 to 1
	Default 1

ethernet

Syntax	ethernet
Context	config>oam-pm>session
Description	This command allows the operator to enter the hierarchy to configure the Ethernet specific source and destination information, the priority, and the Ethernet tests tools on the launch point.

dest-mac

Syntax	dest-mac <i>ieee-address</i> no dest-mac
Context	config>oam-pm>session>ethernet
Description	This command defines the destination MAC address of the peer MEP and sets the destination MAC address in the layer two header to match. This must be a unicast address. The no form of the command removes session parameter.

Parameters	<i>ieee-address</i> — Specifies the layer two 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)

priority

Syntax	priority <i>priority</i>
Context	config>oam-pm>session>ethernet
Description	This command defines 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. The no form of the command removes changes the priority to the default value.
Parameters	<i>priority</i> — Specifies the CoS value.
Values	0 to 7
Default	0

source

Syntax	source mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> no source
Context	config>oam-pm>session>ethernet
Description	This command defines the source launch point identification 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 frames and received as seen under the “show oam-pm statistics session session-name ...” will be zero. The no form of the command removes this session parameter.
Parameters	mep <i>mep-id</i> — Specifies the maintenance association end point identifier of the launch point. Values 1 to 8191 domain <i>md-index</i> — Specifies the maintenance domain (MD) index value of the launch point. Values 1 to 4294967295 association <i>ma-index</i> — Specifies the maintenance association (MA) index value of the launch point. Values 1 to 4294967295

slm

Syntax	slm [test-id test-id] create no slm
Context	config>oam-pm>session>ethernet
Description	This command defines the test-id to be assigned to the synthetic loss test and creates the container to allow the individual test parameters to be configured. The no form of the command removes the SLM 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 an ETH-SLM PDU. Values 0 to 2,147,483,647 create — Keyword to create the test.

dmm

Syntax	dmm [test-id test-id] create no dmm
Context	config>oam-pm>session>ethernet
Description	This command defines 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 an ETH-DMM PDU. Values 0 to 2,147,483,647 create — Keyword to create the test.

lmm

Syntax	lmm [test-id test-id] create no lmm
Context	config>oam-pm>session>ethernet
Description	This command defines the test-id to be assigned to the Tx and Rx counter-based loss test and creates the individual test. LMM does not carry this test-id in the PDU; the value is of local significance. The no form of the command removes the LMM test function from the PM Session.

Parameters *test-id* — Specifies the value to be placed in the 4-byte test id field of an ETH-DMM PDU.

Values 0 to 2,147,483,647

create — Keyword to create the test.

availability

Syntax **availability**

Context config>oam-pm>session>ethernet>Imm

Description This command enables the context to activate, collect, and record availability statistics for LMM tests. These computations are not enabled by default. In order to modify parameters within a session, including these availability parameters, the LMM test must be shut down.

shutdown

Syntax [no] **shutdown**

Context config>oam-pm>session>ethernet>Imm>availability

Description This command deactivates computation and reporting of availability metrics for LMM tests.

 The **no** form of the command activates computation and reporting of availability metrics for LMM tests.

Default shutdown

timing

Syntax **timing frames-per-delta-t frames consec-delta-t deltas chli-threshold threshold**
no timing

Context config>oam-pm>session>ethernet>Imm>availability

Description This command defines various availability parameters for LMM availability testing. This command does not define the probe interval. Validation occurs when the LMM test is activated using the **no shutdown** command. The maximum size of the availability window cannot exceed 100 s (100 000 ms). LMM test activation fails if the availability window exceeds the maximum value.

 The **no** form of the command restores the default values for all timing parameters, and uses those values to compute availability and set the loss frequency.

Parameters	frames-per-delta-t — defines the size of the small measurement window. Each delta-t will be marked as a high-loss interval or non-high-loss interval based on the flr-threshold. The size of the delta-t measurement is the product of the number of frames and the interval. <i>frames</i> — the number of SLM frames that define the size of the delta-t Values 1 to 50 Default 10 consec-delta-t — the number of consecutive delta-t small measurement intervals that make up the sliding window over which availability and unavailability will be determined. Transitions from one state to another will occur when the consec-delta-t are in a new state. The sliding window cannot exceed 100 s. <i>deltas</i> — the number of consecutive delta-t measurement intervals used for the sliding window Values 2 to 10 Default 10 chli-threshold — the number of consecutive high-loss intervals (unavailable delta-t) that, when reached or exceeded, will increment the CHLI counter. 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. <i>threshold</i> — the number of consecutive unavailable delta-t intervals that will cause the CHLI counter to be incremented Values 1 to 9 Default 5

data-tlv-size

Syntax	data-tlv-size <i>octets</i> no data-tlv-size
Context	config>oam-pm>session>ethernet>slm config>oam-pm>session>ethernet>dmm
Description	This command allows the operator to add an optional Data TLV to PDU and increase the frame on the wire by the specified amount. This value is not the size of the frame on the wire. It is the size of the addition padding added to the PDU. The no form of the command removes the optional TVL.
Parameters	<i>octets</i> — Octet size of the optional Data TLV. Values 0 3 to 2000 Default 0

shutdown

Syntax	[no] shutdown
Context	config>oam-pm>session>ethernet>slm config>oam-pm>session>ethernet>dmm config>oam-pm>session>ethernet>Imm
Description	This command activates and deactivates the individual test. When the test is shutdown, no active measurements are being 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 individual test.
Default	shutdown

test-duration

Syntax	test-duration <i>seconds</i> no test-duration
Context	config>oam-pm>session>ethernet>slm config>oam-pm>session>ethernet>dmm config>oam-pm>session>ethernet>Imm
Description	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. On-demand tests do not start until the config>oam-pm>session>start command has been issued and they will stop when the config>oam-pm>session>stop command is issued. The no form of the command will remove a previously configured test-duration and allow the test to execute until manually stopped.
Default	no test-duration
Parameters	<i>seconds</i> — The number of seconds the test will execute from its start time. Values 1 to 86400

flr-threshold

Syntax	flr-threshold <i>percentage</i> no flr-threshold
Context	config>oam-pm>session>ethernet>slm config>oam-pm>session>ethernet>Imm>availability

Description This command defines the frame loss threshold used to determine if the delta-t is available or unavailable. An individual delta-t with a frame loss threshold equal to 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.

The **no** form of the command restores the default value of 50%.

Parameters *percentage* — The percentage of the threshold.

Values 1 to 100

Default 50 percent

timing

Syntax **timing frames-per-delta-t** *frames* **consec-delta-t** *deltas* **interval** *milliseconds* **chli-threshold** *threshold*
no timing

Context config>oam-pm>session>ethernet>slm

Description This command defines various availability parameters and the probe spacing (interval) for the SLM frames. The maximum size of the availability window cannot exceed 10s (10,000ms).

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 in “no shutdown” it will always have timing parameters, default or operator configured.

Parameters **frames-per-delta-t** — Defines the size of the small measurement window. 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.

frames — The number of SLM frames that define the size of the delta-t.

Values 1 to 50

Default 10

consec-delta-t — The number of consecutive delta-t small measurement intervals that make up the sliding window over which availability and unavailability will be determined. Transitions from one state to another will occur when the consec-delta-t are now in a new state.

deltas — The number of consecutive delta-t used for the sliding window.

Values 2 to 10

Default 10

interval — The message period, or probe spacing, for the transmission of the SLM frame.

milliseconds — The number of milliseconds between the transmission of the SLM frames. The default value for the SLM interval is different than the default interval for DMM. This is intentional

Values 100 | 1000

Default 100

chli-threshold — Number of consecutive high loss intervals (unavailable delta-t) that when equal to or exceeded will increment the CHLI counter. A CHLI counter is an indication that the sliding window is available but has crossed a threshold consecutive of unavailable delta-t intervals. A CHLI can only be incremented once during a sliding window and will only be incremented during times of availability.

threshold — The number of consecutive unavailable delta-t that will cause the CHLI counter to be incremented

Values 1 to 9

Default 5

interval

Syntax	interval <i>milliseconds</i> no interval
Context	config>oam-pm>session>ethernet>dmm config>oam-pm>session>ethernet>lmm
Description	This command defines the message period or probe spacing for the transmission of the DMM or LMM frame. The no form of the command sets the interval to the default. If an LMM test is in no shutdown it will always have timing parameters, whether default or operator configured.
Parameters	<i>milliseconds</i> — The number of milliseconds between the transmission of the DMM or LMM frames. The default value for the DMM or LMM interval is different than the default interval for SLM. This is intentional. Values 100 1000 10000 Default 1000

loss-events

Syntax	loss-events
Context	config>oam-pm>session>ethernet>slm config>oam-pm>session>ethernet>lmm config>oam-pm>session>ip>twamp-light

Description This context allows the operator to define the loss events and thresholds that are to be tracked.

avg-flr-event

Syntax	avg-flr-event { forward backward } threshold <i>raise-threshold-percent</i> [clear <i>clear-threshold-percent</i>] [no] avg-flr-event { forward backward }
Context	config>oam-pm>session>ethernet>slm>loss-events config>oam-pm>session>ethernet>lmm>loss-events config>oam-pm>session>ip>twamp-light>loss-events
Description	This command sets the frame loss ratio threshold configuration that will be applied and checked at the end of the measurement interval for the specified direction. This is a percentage based on average frame loss ratio over the entire measurement interval. If the [clear <i>clear-threshold-percent</i>] is not specified the traffic crossing alarm will be stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. Each unique event can only be raised once within measurement interval. If the optional clear threshold is specified the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another will not be raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event will be raised under that condition. The no version of this command removes the event threshold for frame loss ratio. The direction must be included with the no command.
Default	no avg-flr-event forward no avg-flr-event backward
Parameters	forward — The threshold is applied to the forward direction value. backward — The threshold is applied to the backward direction value. threshold <i>raise-threshold-percent</i> — The rising percentage that determines when the event is to be generated. Values 0.001 to 100.000 clear <i>clear-threshold-percent</i> — An optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated. Values 0.000 to 99.999 A value 0.000 means there FLR must be 0.000.

chli-event

Syntax	chli-event {forward backward aggregate} threshold <i>raise-threshold</i> [clear <i>clear-threshold</i>] [no] chli-event {forward backward aggregate}
Context	config>oam-pm>session>ethernet>slm>loss-events config>oam-pm>session>ethernet>lmm>loss-events config>oam-pm>session>ip>twamp-light>loss-events
Description	This command sets the consecutive high loss interval (CHLI) threshold to be monitored and the associated thresholds using the counter of the specified direction. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the [clear <i>clear-threshold</i>] is not specified the traffic crossing alarm will be stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. Each unique event can only be raised once within measurement interval. If the optional clear threshold is specified the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another will not be raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event will be raised under that condition. The no version of this command removes the event threshold for frame loss ratio. The direction must be included with the no command.
Default	no chli-event forward no chli-event backward no chli-event aggregate
Parameters	forward — The threshold is applied to the forward direction count. backward — The threshold is applied to the backward direction count. aggregate — The threshold is applied to the aggregate count (sum of forward and backward). threshold <i>raise-threshold</i> — A numerical value compared to the CHLI counter that is the rising threshold that determines when the event is to be generated, when the percentage of loss value is reached. Values 1 to 864000 clear <i>clear-threshold</i> — An optional numerical value compared to the CHLI counter used for stateful behavior that allows the operator to configure a value lower than the rising percentage to indicate when the clear event should be generated. Values 0 to 863999 A value of zero means the CHLI counter must be 0.

hli-event

Syntax	hli-event { forward backward aggregate } threshold <i>raise-threshold</i> [clear <i>clear-threshold</i>] [no] hli-event { forward backward aggregate }
Context	config>oam-pm>session>ethernet>slm>loss-events config>oam-pm>session>ethernet>lmm>loss-events config>oam-pm>session>ip>twamp-light>loss-events
Description	This command sets the high loss interval (HLI) threshold to be monitored and the associated thresholds using the counter of the specified direction. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the [clear <i>clear-threshold</i>] is not specified the traffic crossing alarm will be stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. Each unique event can only be raised once within measurement interval. If the optional clear threshold is specified the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another will not be raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event will be raised under that condition. The no version of this command removes the event threshold for frame loss ratio. The direction must be included with the no command.
Default	no hli-event backward no hli-event aggregate
Parameters	forward — The threshold is applied to the forward direction count. backward — The threshold is applied to the backward direction count. aggregate — The threshold is applied to the aggregate count (sum of forward and backward). threshold <i>raise-threshold</i> — The rising threshold that determines when the event is to be generated, when the percentage of loss value is reached. Values 1 to 864000 clear <i>clear-threshold</i> — An optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated. Values 0 to 863999 A value of zero means the HLI counter must be 0.

unavailability-event

Syntax **unavailability-event** {**forward** | **backward** | **aggregate**} **threshold** *raise-threshold* [**clear**

	<code>clear-threshold]</code> <code>[no] unavailability-event {forward backward aggregate}</code>
Context	config>oam-pm>session>ethernet>slm>loss-events config>oam-pm>session>ethernet>lmm>loss-events config>oam-pm>session>ip>twamp-light>loss-events
Description	This command sets the threshold to be applied to the overall count of the unavailability indicators, not transitions, per configured direction. This value is compared to the 32 bit unavailability counter specific to the direction which tracks the number of individual delta-ts that have been recorded as unavailable. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the <code>[clear clear-threshold]</code> is not specified, the traffic crossing alarm will be stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. Each unique event can only be raised once within measurement interval. If the optional clear threshold is specified the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised, another will not be raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event will be raised under that condition. The no version of this command removes the event threshold for frame loss ratio. The direction must be included with the no command.
Default	no unavailable-event forward no unavailable-event backward no unavailable-event aggregate
Parameters	forward — The threshold is applied to the forward direction count. backward — The threshold is applied to the backward direction count. aggregate — The threshold is applied to the aggregate count (sum of forward and backward). threshold — A numerical value compared to the unavailability counter that is the rising threshold that determines when the event is to be generated, when value reached. Values 1 to 864000 clear clear-threshold — An optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated. Values 0 to 863999 A value of zero means the unavailability counter must be 0

undet-availability-event

Syntax	undet-availability-event { forward backward aggregate } threshold <i>raise-threshold</i> [clear <i>clear-threshold</i>] [no] undet-availability-event { forward backward aggregate }
Context	config>oam-pm>session>ethernet>slm>loss-events config>oam-pm>session>ethernet>lmm>loss-events config>oam-pm>session>ip>twamp-light>loss-events
Description	This command sets the threshold to be applied to the overall count of the undetermined availability indicators, not transitions, per configured direction. This value is compared to the 32 bit unavailability counter specific to the direction which tracks the number of individual delta-ts that have been recorded as undetermined available. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the [clear <i>clear-threshold</i>] is not specified the traffic crossing alarm will be stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. Each unique event can only be raised once within measurement interval. If the optional clear threshold is specified the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another will not be raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event will be raised under that condition. The no version of this command removes the event threshold for frame loss ratio. The direction must be included with the no command.
Default	no undetermined-available-event forward no undetermined-available-event backward no undetermined-available-event aggregate
Parameters	forward — The threshold is applied to the forward direction count. backward — The threshold is applied to the backward direction count. aggregate — The threshold is applied to the aggregate count (sum of forward and backward). threshold <i>raise-threshold</i> — The rising threshold that determines when the event is to be generated, when value reached. Values 1 to 864000 clear <i>clear-threshold</i> — An optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated. Values 0 to 863999 A value of zero means the undetermined availability counter must be 0.

undet-unavailability-event

Syntax	undet-availability-event { forward backward aggregate } threshold <i>raise-threshold</i> [clear <i>clear-threshold</i>] [no] undet-availability-event { forward backward aggregate }
Context	config>oam-pm>session>ethernet>slm>loss-events config>oam-pm>session>ethernet>lmm>loss-events config>oam-pm>session>ip>twamp-light>loss-events
Description	This command sets the threshold to be applied to the overall count of the undetermined unavailability indicators, not transitions, per configured direction. This value is compared to the 32 bit unavailability counter specific to the direction which tracks the number of individual delta-ts that have been recorded as undetermined unavailable. The aggregate is a function of summing forward and backward. This value is only used as a threshold mechanism and is not part of the stored statistics. If the [clear <i>clear-threshold</i>] is not specified the traffic crossing alarm will be stateless. Stateless means the state is not carried forward to other measurement intervals. Each measurement interval is analyzed independently and without regard to any previous window. Each unique event can only be raised once within measurement interval. If the optional clear threshold is specified the traffic crossing alarm uses stateful behavior. Stateful means each unique previous event state is carried forward to following measurement intervals. If a threshold crossing event is raised another will not be raised until a measurement interval completes and the clear threshold has not been exceeded. A clear event will be raised under that condition. The no version of this command removes the event threshold for frame loss ratio. The direction must be included with the no command.
Default	no undet-unavailable-event forward no undet-unavailable-event backward no undet-unavailable-event aggregate
Parameters	forward — The threshold is applied to the forward direction count. backward — The threshold is applied to the backward direction count. aggregate — The threshold is applied to the aggregate count (sum of forward and backward). threshold <i>raise-threshold</i> — The rising threshold that determines when the event is to be generated, when value reached. Values 1 to 864000 clear <i>clear-threshold</i> — An optional value used for stateful behavior that allows the operator to configure a percentage of loss value lower than the rising percentage to indicate when the clear event should be generated. Values 0 to 863999 A value of zero means the undetermined availability counter must be 0.

3.12.2.1.12 LDP TreeTrace Commands

ldp-treetrace

Syntax	ldp-treetrace { prefix <i>ip-prefix/mask</i> } [max-ttl <i>ttl-value</i>] [max-path <i>max-paths</i>] [timeout <i>timeout</i>] [retry-count <i>retry-count</i>] [fc <i>fc-name</i> [profile <i>profile</i>]] [downstream-map-tlv { dsmap ddmap }]
Context	oam
Description	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 <i>ip-prefix/mask</i> — Specifies the address prefix and subnet mask of the target BGP IPv4 label route. max-ttl <i>max-label-ttl</i> — The maximum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer. Values 1 to 255 Default 30 max-paths <i>max-paths</i> — The maximum number of paths for a ldp-treetrace test, expressed as a decimal integer. Values 1 to 255 Default 128 timeout <i>timeout</i> — 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 to 60 Default 3 fc <i>fc-name</i> — The fc and profile parameters are used to indicate the forwarding class and profile 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 and profile 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 and profile 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 22](#) summarizes this behavior.

Table 22 **Idp-treetrace Request Packet and Behavior**

cpm (sender node)	echo request packet: • packet{tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
outgoing interface (sender node)	echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	echo request packet: • packet{tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
cpm (responder node)	echo reply packet: • packet{tos=1, fc2, profile2}
outgoing interface (responder node)	echo reply packet: • pkt queued as {fc2, profile2} • 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)	echo reply packet: • packet{tos1, exp2} • exp2 mapped to {fc1, profile1} as per classification in network QoS policy of incoming interface

Values be, l2, af, l1, 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 to 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**}.

Output

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	This command enables the context to configure Operations, Administration, and Maintenance test parameters.

ldp-treetrace

Syntax	[no] ldp-treetrace
Context	config>test-oam
Description	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.

Output

Sample Output Over a Numbered IP Interface

```
*A:Dut-B# oam ldp-treetrace prefix 10.20.1.5/32

ldp-treetrace 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
```

Sample output Over an Unnumbered IP Interface

```
*A:Dut-A# oam ldp-treetrace prefix 10.20.1.6/32 downstream-map-tlv dsmap

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
```

fc

Syntax	fc <i>fc-name</i> [profile {in out}] no fc
Context	config>test-oam>ldp-treetrace
Description	This command indicates the forwarding class and profile 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 and profile 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 and profile 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 23](#) summarizes this behavior.

Table 23 fc Request Packet and Behavior

cpm (sender node)	echo request packet: • packet{tos=1, fc1, profile1} • fc1 and profile1 are as entered by user in OAM command or default values • tos1 as per mapping of {fc1, profile1} to IP precedence in network egress QoS policy of outgoing interface
outgoing interface (sender node)	echo request packet: • pkt queued as {fc1, profile1} • ToS field=tos1 not remarked • EXP=exp1, as per mapping of {fc1, profile1} to EXP in network egress QoS policy of outgoing interface
Incoming interface (responder node)	echo request packet: • packet{tos1, exp1} • exp1 mapped to {fc2, profile2} as per classification in network QoS policy of incoming interface
cpm (responder node)	echo reply packet: • packet{tos=1, fc2, profile2}

Table 23 fc Request Packet and Behavior (Continued)

outgoing interface (responder node)	echo reply packet: • pkt queued as {fc2, profile2} • 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)	echo reply packet: • 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, l1, h2, ef, h1, nc

profile {in | out} — Specifies the profile value to be used with the forwarding class specified in the **fc-name** parameter.

path-discovery

Syntax path-discovery

Context config>test-oam>ldp-treetrace

Description This command creates the context to configure the LDP ECMP OAM path discovery.

The ingress LER builds the ECM 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 learned and added into the FEC database. The maximum number of FECs to be discovered with the tree building feature is limited to 500. 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 <i>minutes</i> no interval
Context	config>test-oam>ldp-treetrace>path-discovery
Description	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	<i>minutes</i> — Specifies the number of minutes to wait before repeating the LDP tree auto discovery process. Values 60 to 1440

max-path

Syntax	max-path <i>max-paths</i>
Context	config>test-oam>ldp-treetrace>path-discovery
Description	This command configures the maximum number of ECMP paths the path discovery will attempt to discover for each run every interval minutes. The no option resets the timeout to its default value.
Default	128
Parameters	<i>max-paths</i> — Specifies the tree discovery maximum path. Values 1 to 128

max-ttl

Syntax	max-ttl <i>ttl-value</i>

Context	config>test-oam>ldp-treetrace>path-discovery
Description	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	<i>tvl-value</i> — Specifies the maximum label time-to-live value for an LSP trace request during the tree discovery. Values 1 to 255

policy-statement

Syntax	policy-statement <i>policy-name</i> [...(up to 5 max)]
Context	config>test-oam>ldp-treetrace>path-discovery
Description	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 learned 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	<i>policy-name</i> — 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 <i>retry-count</i> no retry-count
Context	config>oam-test>ldp-treetrace>path-discovery

```
config>oam-test>ldp-treetrace>path-probing
```

Description	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	<i>retry-count</i> — Specifies the maximum number of consecutive timeouts allowed before failing a path probe (ping). Values 1 to 10

timeout

Syntax	timeout <i>timeout</i> no timeout
Context	config>test-oam>ldp-treetrace>path-discovery
Description	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.
Default	30
Parameters	<i>timeout</i> — Specifies the timeout parameter, in seconds, within a range of 1 to 60, expressed as a decimal integer. Values 1 to 60

path-probing

Syntax	path-probing
Context	config>test-oam>ldp-treetrace
Description	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 <i>minutes</i> no interval
Context	config>test-oam>ldp-treetrace>path-probing
Description	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 the LDP tree trace path discovery. The no option resets the interval to its default value.
Default	1
Parameters	<i>minutes</i> — Specifies the number of minutes to probe all active ECMP paths for each LDP FEC. Values 1 to 60

timeout

Syntax	timeout <i>timeout</i> no timeout
Context	config>test-oam>ldp-treetrace>path-probing
Description	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. The no option resets the timeout to its default value.
Default	1

Parameters *timeout* — Specifies the timeout parameter, in minutes, with a range of 1 to 3 minutes, expressed as a decimal integer.

Values 1 to 3

mpls-time-stamp-format

Syntax **mpls-time-stamp-format {rfc4379 | unix}**

Context config>test-oam

Description This command configures the format of the timestamp used by for lsp-ping, lsp-trace, p2mp-lsp-ping and p2mp-lsp-trace, vccv-ping, vccv-trace, and lsp-trace.

If **rfc4379** is selected, then the timestamp is in seconds and microseconds since 1900, otherwise it is in seconds and microseconds since 1970.

Changing this system-wide setting does not affect tests that are currently in progress, but SAAs will start to use the new timestamp when they are restarted. When an SR OS receives an echo request, it will reply with the locally configured timestamp format, and will not try to match the timestamp format of the incoming echo request message.

Default unix

Parameters **rfc4379** — Specifies the RFC 4379 time stamp format. The timestamp's **seconds** field holds the integral number of seconds since 1-Jan-1900 00:00:00 UTC. The timestamp's **microseconds** field contains a microseconds value in the range 0 to 999999. This setting is used to inter-operate with network elements which are fully compliant with RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*, (such as an SR OS system with the same setting, or any other RFC 4379 compliant router).

unix — Specifies the Unix time stamp format. The time stamps **seconds** field holds a Unix time, the integral number of seconds since 1-Jan-1970 00:00:00 UTC. The time stamps **microseconds** field contains a microseconds value in the range 0 to 999999. This setting is used to inter-operate with network elements which send and expect a 1970-based timestamp in MPLS Echo Request/Reply PDUs (such as an SR OS system with the same setting, or an SR OS system running software earlier than R8.0 R4).

mpls-echo-request-downstream-map

Syntax **mpls-echo-request-downstream-map {dsmap | ddmap}**
no mpls-echo-request-downstream-map

Context config>test-oam

Description 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.

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.

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 | ddmmap | 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 | ddmmap | 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 DMAP or DMAP TLV in an lsp-ping message.

In addition to performing the same features as the DMAP TLV, the new DMAP 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. This includes the case of explicit configuration of the LDP-BGP stitching in which the BGP label route is active in Route Table Manager (RTM) and the case of a BGP IPv4 label route resolved to the LDP FEC due to the IGP route of the same prefix active in RTM. In this case, 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 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 DMAP 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 DMAP 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 DMAP 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 DMAP TLV in a LSP trace packet and LDP tree trace packet.

Default dsmap

Output

Sample Output for LDP-over-RSVP




```
Testing LDP FEC of Node C with DMAP TLV
-----
*A:Dut-A#
```

```
*A:Dut-A# oam lsp-trace prefix 10.20.1.3/32 downstream-map-tlv dsmmap 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)
       label[2]=131071 protocol=3(LDP)
       fecchange[1]=PUSH fecctype=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#

```

Sample Output for LDP-BGP Stitching



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 ddmmap 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#

```



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 ddmmap 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#

```

3.12.2.1.13 TWAMP Commands

twamp

Syntax	twamp
Context	config>test-oam
Description	This command enables TWAMP functionality.
Default	TWAMP is disabled.

server

Syntax	retry-count <i>retry-count</i>
Context	config>test-oam>twamp
Description	This command configures the node for TWAMP server functionality.
Default	TWAMP is disabled.

prefix

Syntax	prefix <i>address/prefix-length</i> [create] no prefix <i>address/prefix-length</i>
Context	config>test-oam>twamp>server
Description	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

Parameters	<i>address</i> — An IPv4 or IPv6 address prefix (with host bits set to 0).
	<i>prefix length</i> — The prefix length.
Values	0 to 128
Default	none

description

Syntax	description <i>text</i> no description
Context	config>test-oam>twamp>server>prefix
Description	Use this command to configure a description for the TWAMP server prefix table. The no form of the command removes the configuration.
Default	no description
Parameters	<i>text</i> — The TWAMP server description, up to 80 characters in length.

max-conn-prefix

Syntax	max-conn-prefix <i>count</i> no max-conn-prefix
Context	config>test-oam>twamp>server>prefix
Description	This command configures the maximum number of 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 (32).
Default	no max-conn-prefix
Parameters	<i>count</i> — The maximum number of control connections.
Values	0 to 64
Default	32

max-conn-server

Syntax	max-conn-server <i>count</i> no max-conn-server

Context	config>test-oam>twamp>server				
Description	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 (32).				
Default	no max-conn-server				
Parameters	<i>count</i> — The maximum number of control connections. <table><tr><td>Values</td><td>0 to 64</td></tr><tr><td>Default</td><td>32</td></tr></table>	Values	0 to 64	Default	32
Values	0 to 64				
Default	32				

inactivity-timeout

Syntax	inactivity-timeout <i>seconds</i> no inactivity-timeout				
Context	config>test-oam>twamp>server				
Description	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 in-progress tests are terminated. The no form of the command sets the default value (1800 s.)				
Default	no inactivity-timeout				
Parameters	<i>seconds</i> — The duration of the inactivity timeout. <table><tr><td>Values</td><td>0 to 3600</td></tr><tr><td>Default</td><td>1800</td></tr></table>	Values	0 to 3600	Default	1800
Values	0 to 3600				
Default	1800				

max-sess-prefix

Syntax	max-sess-prefix <i>count</i> no max-sess-prefix
Context	config>test-oam>twamp>server>prefix
Description	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 sets the default value (32).

Default	no max-sess-prefix
Parameters	<i>count</i> — The maximum number of concurrent test sessions.
Values	0 to 128
Default	32

max-sess-server

Syntax	max-sess-server <i>count</i> no max-sess-server
Context	config>test-oam>twamp>server
Description	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 means to go with a default value of 32.
Default	no max-sessions
Parameters	<i>count</i> — The maximum number of concurrent test sessions.
Values	0 to 128
Default	32

3.12.2.1.14 TWAMP Light Commands

twamp-light

Syntax	twamp-light
Context	config>router config>service>vpn config>test-oam>twamp
Description	This command enables the context for configuring TWAMP Light parameters.

inactivity-timeout

Syntax	inactivity-timeout <i>time</i> no inactivity-timeout

Context	config>test-oam>twamp>twamp-light				
Description	This command configures the length of time to maintain stale state on the session reflector. Stale state is test data that has not been refreshed or updated by newly arriving probes for that specific test in a predetermined length of time. Any single reflector can maintain up state for a maximum of 12,000 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 sets the default value of 100.				
Parameters	<i>time</i> — The value in seconds for maintaining stale state. <table> <tr> <td>Values</td><td>10 to 100</td></tr> <tr> <td>Default</td><td>100</td></tr> </table>	Values	10 to 100	Default	100
Values	10 to 100				
Default	100				

reflector

Syntax	reflector [udp-port <i>udp-port-number</i>] [create] no reflector		
Context	config>router>twamp-light config>service>vpn>twamp-light		
Description	Use this command to configure TWAMP Light session reflector parameters and to enable TWAMP Light functionality with the no shutdown command. The udp-port keyword and value must be specified with the create keyword. An error message is generated if the specific UDP port is unavailable.		
Parameters	<i>udp-port-number</i> — Specifies the UDP port number. A strictly enforced restricted range has been introduced. The TWAMP Light session reflector must be brought in line with this new restriction prior upgrading or rebooting from any previous release if there is an active TWAMP Light session reflector configured. Failure to do so will prevent an ISSU operation from proceeding and will fail to activate any reflector outside of the enforced range. Note that in the Two-Way Active Measurement Protocol Light (TWAMP Light) section for a complete description. This parameter is required and specifies the destination udp-port that the session reflector will use to listen for TWAMP Light packets. The session controller launching the TWAMP Light packets must be configured with the same destination UDP port as part of the TWAMP Light test. The IES service will use the destination UDP port that is configured under the router context. Only one udp-port may be configured per unique context. <table> <tr> <td>Values</td><td>64364 to 64373</td></tr> </table>	Values	64364 to 64373
Values	64364 to 64373		

prefix

Syntax	prefix <i>ip-prefix/prefix-length</i> [create] no prefix

Context	config>router>twamp-light>reflector config>service>vprn>twamp-light>reflector										
Description	Use this command to define which TWAMP Light packet prefixes the reflector will process. The no form of the command with the specific prefix removes the accepted source.										
Parameters	create — Instantiates the prefix list <i>ip-prefix/prefix-length</i> — The IPv4 or IPv6 address and length (applies t the 7750 SR and 7950 XRS). Values <table><tr><td>ipv4-prefix:</td><td>a.b.c.d (host bits must be 0)</td></tr><tr><td>ipv4-prefix-le:</td><td>0 to 32</td></tr><tr><td>ipv6-prefix:</td><td>x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D</td></tr><tr><td>ipv6-prefix-le:</td><td>0 to 128</td></tr><tr><td>ipv6-address:</td><td>x:x:x:x:x:x:x x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D</td></tr></table>	ipv4-prefix:	a.b.c.d (host bits must be 0)	ipv4-prefix-le:	0 to 32	ipv6-prefix:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D	ipv6-prefix-le:	0 to 128	ipv6-address:	x:x:x:x:x:x:x x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D
ipv4-prefix:	a.b.c.d (host bits must be 0)										
ipv4-prefix-le:	0 to 32										
ipv6-prefix:	x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D										
ipv6-prefix-le:	0 to 128										
ipv6-address:	x:x:x:x:x:x:x x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D										

description

Syntax	description <i>description-string</i> no description
Context	config>router>twamp-light>reflector>prefix config>service>vprn>twamp-light>reflector>prefix config>router>twamp-light>reflector config>service>vprn>twamp-light>reflector
Description	Use this command to configure a text description that gets 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	<i>description-string</i> — The description character string. Allowed values are any characters up to 80 characters in length, composed of printable, 7-bit ASCII characters. If the string contains special characters (for example, #, \$, or spaces), the entire string must be enclosed in double quotes

shutdown

Syntax	shutdown no shutdown
Context	config>router>twamp-light>reflector config>service>vprn>twamp-light>reflector
Description	Use this command to disable or enable TWAMP Light functionality within the context where the configuration exists, either the base router instance or the service. Enabling the base router context enables the IES prefix list since the IES service uses the configuration under the base router instance. The no form of the command allows the router instance or the service to accept TWAMP Light packets for processing.
Default	shutdown

ip

Syntax	ip
Context	config>oam-pm>session>ip
Description	Use this command to enter the context to configure the IP-specific source and destination information, the priority, and the IP test tools on the launch point.

twamp-light

Syntax	twamp-light [test-id <i>test-id</i>] [create] no twamp-light
Context	config>oam-pm>session>ip
Description	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	<i>test-id</i> — Specifies the value of the 4-byte local test identifier not sent in the TWAMP Light packets. Values 0 to 2,147,483,647 create — Keyword to create the test.

source

Syntax	source <i>ip-address</i> no source
Context	config>oam-pm>session>ip
Description	Use this command to define 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	source — Keyword that indicates the launch point. <i>ip-address</i> — This mandatory parameter is required in order to validate the TWAMP Light response received from the reflector. The initial source must be the destination in the response. Values - IPv4 address in the form a.b.c.d - IPv6 address in the form x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D (no multicast addresses)



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

destination

Syntax	destination ip-address no destination
Context	config>oam-pm>session>ip
Description	Use this command to define the destination IP address that will be assigned to the TWAMP Light packets. The destination address must be included in the prefix list on the session reflector within the configured context in order to allow the reflector to process the inbound TWAMP Light packets. The no form of the command removes the destination parameters.

Default	no destination
Parameters	destination — Keyword that indicates the destination of the packet <i>ip-address</i> — Parameter that specifies the IP address of the IP peer to which the packet is directed. Values IPv4 address in the form a.b.c.d Values IPv6 address in the form x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D (no multicast addresses)

dest-udp-port

Syntax	dest-udp-port <i>udp-port-number</i> no dest-udp-port
Context	config>oam-pm>session>ip
Description	Use this command to define the destination UDP port on outbound TWAMP Light packets sent from the session controller. The destination UDP port must match the UDP port value configured on the TWAMP Light reflector that will be responding to this specific TWAMP Light test. The no form of the command removes the destination UDP port setting.
Default	no dest-udp port
Parameters	<i>udp-port-number</i> — The udp source port. Values 1 to 65535

source-udp-port

Syntax	source-udp-port <i>udp-port-number</i> no source-udp-port
Context	config>oam-pm>session>ip
Description	Optional command that 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 included the source udp 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 setting when the default allocation is used.

Default dynamic source udp port allocation

Parameters *udp-port-number* — The udp source port.

Values 64374 to 64383

forwarding

Syntax **forwarding** {**next-hop** *ip-address* | **interface** *interface-name* | **bypass-routing**}
no forwarding

Context config>oam-pm>session>ip

Description Use this optional command to influence 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 options and enables the default forwarding logic.

Default no forwarding

Parameters **next-hop** — Specifies the IP next hop on the path.

ip-address — Specifies the 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)

x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D
(no multicast addresses)



Note: ipv6-address applies to the 7750 SR and 7950 XRS only.

interface *interface-name* — Specifies the name used to refer to 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.

bypass-routing — Specifies to send the packet to a host on a directly attached network, bypassing the routing table.

fc

Syntax	fc { be l2 af l1 h2 ef h1 nc } no fc
Context	config>oam-pm>session>ip
Description	Use this command to set the forwarding class designation for TWAMP Light packets that will be 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	be
Parameters	be — Specifies best effort. l2 — Specifies low-2. af — Specifies assured. l1 — Specifies low-1. h2 — Specifies high-2. ef — Specifies expedited. h1 — Specifies high-1. nc — Specifies network control.

profile

Syntax	profile { in out } no profile
Context	config>oam-pm>session>ip
Description	Use this command to define whether the TWAMP Light PDU packet should be treated as in-profile or out-of-profile. The default has been selected because the forwarding class defaults to best effort. The no form of the command restores the default value.
Default	out
Parameters	in — Specifies that the TWAMP Light PDU packet will be sent as in-profile. out — Specifies that the TWAMP Light PDU packet will be sent as out-of-profile.

tth

Syntax	tth <i>time-to-live</i>

no ttl

Context	config>oam-pm>session>ip
Description	Use this command to define the value of the TTL field of the packet header. The no form of the command restores the default value.
Default	225
Parameters	<i>time-to-live</i> — Specifies the value to be used in the TTL field. Values 1 to 255

router

Syntax	router { <i>base</i> <i>routing-instance</i> service-name <i>service-name</i> } no router
Context	config>oam-pm>session>ip
Description	Use this command to define the source context from which the TWAMP Light packet will be launched. The routing instance and service name must be a VPRN instance. The no form of the command restores the default value.
Default	base
Parameters	base — Specifies that the TWAMP Light packet will be launched from the base routing instance. <i>routing-instance</i> — Specifies the service identifier from which the TWAMP Light packet is launched. service-name — Specifies the that the TWAMP Light packet will be launched from a service context. <i>service-name</i> — Specifies the service from which the TWAMP Light packet is launched. Values up to 64 characters in length

pad-size

Syntax	pad-size <i>padding</i> no pad-size
Context	config>oam-pm>session>ip>twamp-light

Description	Use this command to define the amount by which the TWAMP Light packet will be 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, the pad size must be configured to a minimum of 27 bytes. The no form of the command removes all padding.
Default	0
Parameters	<i>padding</i> — Specifies the value, in octets, to pad the TWAMP Light packet. Values 0 to 2000

record-stats

Syntax	record-stats {delay loss delay-and-loss} [no] record-stats
Context	config>oam-pm>session>ip>twamp-light
Description	This option provides the ability to determine which statistics are recorded. The TWAMP-Light PDU can report on both delay and loss using a single packet. The operator may choose which statistics they would like to report. Only delay recording is on by default. All other metrics are ignored. In order to change what is being recorded and reported, the TWAMP-Light session must be shutdown. This is required because the single packet approach means the base statistics are shared between the various datasets. Issuing a “no shutdown” will clear previous all non-volatile memory for the session and allocate new memory blocks. All the parameters under this context are mutually exclusive. The no version of the command restores the default “delay” only.
Default	delay
Parameters	delay — Delay only recording (the default). loss — Loss only recording. delay-and-loss — Delay and loss reporting.

loss

Syntax	loss
Context	config>oam-pm>session>ip>twamp-light
Description	This command configures loss parameters for the TWAMP-Light test.

fir-threshold

Syntax	[no] fir-threshold <i>percentage</i>
Context	config>oam-pm>session>ip>twamp-light>loss
Description	This command defines the frame loss threshold 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. The no form of the command restores the default value of 50%.
Parameters	<i>percentage</i> — The percentage of the threshold. Values 0 to 100 Default 50 percent

timing

Syntax	[no] timing frames-per-delta-t <i>frames</i> consec-delta-t <i>deltas</i> chli-threshold <i>threshold</i>
Context	config>oam-pm>session>ip>twamp-light>loss
Description	This command defines various availability parameters but not the probe interval. A single TWAMP-Light frame is used to collect both delay and loss metrics the interval is common to both and as such not unique per metric type. Any TWAMP light test that is attempting to become active will validate the configuration of the timing parameter regardless of which statistics are being recorded. The no form of the command will restore the default values for all timing parameters and use those values to compute availability and set the loss frequency.
Parameters	frames-per-delta-t <i>frames</i> — Defines the size of the small measurement window. Each delta-t will be marked as available or unavailable based on the fir-threshold. The size of the delta-t measurement is the product of the number of frames and the interval. This value defaults to a different value than single probe per metric approaches. Values 1 to 50 Default 1 consec-delta-t <i>deltas</i> — The number of consecutive delta-t small measurement intervals that make up the sliding window over which availability and unavailability will be determined. Transitions from one state to another will occur when the consec-delta-t are now in a new state. The sliding window cannot exceed 100s. Values 2 to 10 Default 10

chli-threshold *threshold* — Number of consecutive high loss intervals (unavailable delta-t) that when equal to or exceeded will increment the CHLI counter. A CHLI counter is an indication that the sliding window is available but has crossed a threshold consecutive of 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

Default 5

interval

Syntax **interval** *milliseconds*
no interval

Context config>oam-pm>session>ip>twamp-light

Description	Use this command to define the message period, or probe spacing, for transmitting a TWAMP Light frame.

The **no** form of the command sets the interval to the default value.

Default	1000

Parameters	<i>milliseconds</i> — Specifies the number of milliseconds between TWAMP Light frame transmission.

Values 100 | 1000 | 10000

test-duration

Syntax **test-duration** *seconds*
 no test-duration

Context config>oam-pm>session>ip>twamp-light

Description	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. On-demand tests do not start until the config>oam-pm>session>start command has been issued and they will stop when the config>oam-pm>session>stop command is issued.

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 0

Parameters	<i>seconds</i> — Specifies the length of time, in seconds, that the TWAMP Light test will run.
Values	1 to 86400

shutdown

Syntax	[no] shutdown
Context	config>oam-pm>session>ip>twamp-light
Description	Use this command to stop a TWAMP Light test. The no form of the command starts a TWAMP Light test.
Default	shutdown

3.12.2.2 Show Commands

The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

saa

Syntax	saa [<i>test-name</i>] [owner <i>test-owner</i>]
Context	show>saa
Description	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	<i>test-name</i> — The optional parameter is used to 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. owner <i>test-owner</i> — 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	

Sample Output

```

*A:bksim130>config>saa>test>trap-gen# show saa mySaaPingTest1
=====
SAA Test Information
=====
Test name           : mySaaPingTest1
Owner name          : TiMOS CLI
Description          : N/A
Accounting policy    : None
Administrative status : Disabled
Test type           : icmp-ping 11.22.33.44
Trap generation      : probe-fail-enable probe-fail-threshold 3
                     : test-fail-enable test-fail-threshold 2
                     : test-completion-enable

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
          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
=====

*A:bksim130>config>saa>test>trap-gen#

*A:bksim130>config>saa>test>trap-gen$ show saa mySaaTraceRouteTest1
=====
SAA Test Information
=====
Test name           : mySaaTraceRouteTest1
Owner name          : TiMOS CLI
Description          : N/A
Accounting policy    : None
Administrative status : Disabled
Test type           : icmp-trace 11.22.33.44
Trap generation      : test-fail-enable test-completion-enable

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
	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

```
=====
*A:bksim130>config>saa>test>trap-gen$
```

```
show saa <test-name>
CFM Loopback:
```

SAA Test Information

```
=====
Test name : CFMLoopbackTest
Owner name      : TiMOS CLI
Description     : N/A
Accounting policy : 1
Continuous      : Yes
Administrative status : Enabled
Test type       : eth-cfm-
loopback 00:01:01:01:01:01 mep 1 domain 1 association 1 interval 1 count 10
Trap generation : None
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
	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: 1

Total number of attempts: 10

Number of requests that failed to be sent out: 0

Number of responses that were received: 10

Number of requests that did not receive any response: 0

Total number of failures: 0, Percentage: 0

(in us)	Min	Max	Average	Jitter
Outbound :	0.000	0.000	0.000	0
Inbound :	0.000	0.000	0.000	0
Roundtrip :	10200	10300	10250	100

Per test packet:

Sequence	Result	Delay(us)
1	Response Received	10300
2	Response Received	10300
3	Response Received	10300
4	Response Received	10200
5	Response Received	10300
6	Response Received	10200
7	Response Received	10300
8	Response Received	10200
9	Response Received	10300
10	Response Received	10300

=====

CFM Traceroute:

=====

SAA Test Information

=====

Test name	: CFMLinkTraceTest
Owner name	: TiMOS CLI
Description	: N/A
Accounting policy	: None
Continuous	: Yes
Administrative status	: Enabled
Test type	: eth-cfm-
linktrace 8A:DB:01:01:00:02	mep 1 domain 1 association 1 interval 1
Trap generation	: None
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
	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: 1
HopIdx: 1
Total number of attempts: 3
Number of requests that failed to be sent out: 0
Number of responses that were received: 3
Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0
(in ms)           Min           Max           Average           Jitter
Outbound  :           0.000         0.000         0.000         0.000
Inbound   :           0.000         0.000         0.000         0.000
Roundtrip :           2.86          3.67          3.15         0.047
Per test packet:
Sequence  Outbound  Inbound  RoundTrip Result
      1      0.000    0.000      3.67 Response Received
      2      0.000    0.000      2.92 Response Received
      3      0.000    0.000      2.86 Response Received

```

```

HopIdx: 2
Total number of attempts: 3
Number of requests that failed to be sent out: 0
Number of responses that were received: 3
Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0
(in ms)           Min           Max           Average           Jitter
Outbound  :           0.000         0.000         0.000         0.000
Inbound   :           0.000         0.000         0.000         0.000
Roundtrip :           4.07          4.13          4.10         0.005
Per test packet:
Sequence  Outbound  Inbound  RoundTrip Result
      1      0.000    0.000      4.10 Response Received
      2      0.000    0.000      4.13 Response Received
      3      0.000    0.000      4.07 Response Received

```

=====

CFM Two Way Delay Measurement:

=====

SAA Test Information

=====

```

Test name           : CFMTwoWayDelayTest
Owner name          : TiMOS CLI
Description         : N/A
Accounting policy   : None
Continuous         : Yes
Administrative status : Enabled
Test type           : eth-cfm-two-way-
delay 00:01:01:01:01:01 mep 1 domain 1 association 1 interval 1
Trap generation     : None
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
           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: 1
HopIdx: 1
Total number of attempts: 3
Number of requests that failed to be sent out: 0
Number of responses that were received: 3
Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0
Total number of failures: 0, Percentage: 0
(in us)      Min      Max      Average      Jitter
Outbound   :      5095      5095      5095      0
Inbound    :      5095      5095      0.000      0
Roundtrip  :     10190     10190     10190      0
Per test packet:
Sequence (in us) Outbound Inbound Delay Delay variation
1         5195      5195    10190      0
2         5195      5195    10190      0
3         5195      5195    10190      0
...
=====
```

Table 24 provides SAA field descriptions.

Table 24 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.

Table 24 SAA Field Descriptions (Continued)

Label	Description (Continued)
Accounting policy	Specifies the associated accounting policy ID.
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	Displays the configured threshold value.
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.

twamp

Syntax `twamp`

Context `show>test-oam`

Description This command enables the context for displaying TWAMP information.

client

Syntax `client {all | ip-address}`

Context	show>test-oam>twamp
Description	This command displays client TWAMP information.
Parameters	all — Displays all client information <i>ip-address</i> — Specifies the IP address of a client
Output	

Sample Output

```

show test-oam twamp client all
=====
Test Session information for Client 6.6.6.6
=====
Index      : 1      State      : Active
SID        : 16 byte hex field
Src Address : SourceIP      Src UDP Port : port
Dst Address : DestIP        Dst UDP Port : port
-----
Index      : 2      State      : Active
SID        : 16 byte hex field
Src Address : SourceIP      Src UDP Port : port
Dst Address : DestIP        Dst UDP Port : port
-----
-----
Number of Sessions: 2
=====
Test Session information for Client 10.10.10.10
=====
Index      : 1      State      : Active
SID        : 16 byte hex field
Src Address : SourceIP      Src UDP Port : port
Dst Address : DestIP        Dst UDP Port : port
-----
Index      : 2      State      : Active
SID        : 16 byte hex field
Src Address : SourceIP      Src UDP Port : port
Dst Address : DestIP        Dst UDP Port : port
-----
-----
Index      : 3      State      : Active
SID        : 16 byte hex field
Src Address : SourceIP      Src UDP Port : port
Dst Address : DestIP        Dst UDP Port : port
-----
-----
Number of Sessions: 3
=====
Test Session information for Client 1234:5678:90ab:cdef:1234:5678:90ab:cdef
=====
Index      : 1      State      : Active
SID        : 16 byte hex field
Src Address : SourceIP      Src UDP Port : port
Dst Address : DestIP        Dst UDP Port : port
-----

```



```
Number of Sessions: 1
=====

show test-oam twamp client 6.6.6.6
=====
Test Session information for Client 6.6.6.6
=====
Index      : 1      State      : Active
SID        : 16 byte hex field
Src Address : SourceIP      Src UDP Port : port
Dst Address : DestIP        Dst UDP Port : port
-----
Index      : 2      State      : Active
SID        : 16 byte hex field
Src Address : SourceIP      Src UDP Port : port
Dst Address : DestIP        Dst UDP Port : port
-----
Number of Sessions: 2
=====
```

server

Syntax	server { all prefix <i>ip-prefix/prefix-length</i> capability }
Context	show>test-oam>twamp
Description	This command displays server TWAMP information.
Parameters	all — Displays all server information prefix — Displays the address prefix of the TWAMP server <i>ip-prefix/prefix-length</i> — Specifies the IP address prefix of the TWAMP server capability — Displays the modes referenced or supported by the TWAMP server, with an RFC reference where those modes are defined.

Output

Sample Output

```
*A:ALA-48# show test-oam twamp server
=====
TWAMP Server (port 862)
=====
Admin State : Up                      Oper State : Up
Up Time     : 0d 00:00:05
Curr Conn   : 1                      Max Conn   : 32
ConnTimeout : 1800                   Conn Reject : 2
Curr Sess   : 2                      Max Sess   : 32
Tests Done  : 5                      Tests Rej   : 0
Tests Abort : 0
TstPktsRx   : 999                   TstPktsTx  : 999
=====
Prefix      : 10.0.0.0/8
```

```

Tests Abort : 0
TstPktsRx   : 999                               TstPktsTx   : 999
=====
Prefix      : 10.0.0.0/8
Description  : NMS-West
=====
Admin State : Up                               Oper State  : Up
Curr Conn   : 1                               Max Conn    : 32
Conn Reject : 0                               Max Sess    : 32
Curr Sess   : 2                               Tests Rej   : 0
Tests Done  : 5                               TstPktsRx  : 999
Tests Abort : 0                               TstPktsTx  : 999
TstPktsRx   : 999
=====
Client      Sessions      Idle    TstPktsRx  TstPktsTx
      Curr/Done/Rej/Abort
-----
10.1.1.1    2/5/0/0        920      999        999
=====
Prefix      : 10.0.0.0/16
Description  : NMS-West-Special
=====
Admin State : Up                               Oper State  : Up
Curr Conn   : 0                               Max Conn    : 32
Conn Reject : 0                               Max Sess    : 32
Curr Sess   : 0                               Tests Rej   : 0
Tests Done  : 0                               TstPktsRx  : 0
Tests Abort : 0                               TstPktsTx  : 0
TstPktsRx   : 0
=====
Client      Sessions      Idle    TstPktsRx  TstPktsTx
      Curr/Done/Rej/Abort
-----
=====

*A:ALA-48# show test-oam twamp server capability
=====
TWAMP Server Supported Modes of Operation with RFC Reference
=====
Bit      Value      Description      RFC
-----
0         1          Unauthenticated  5357
4         16          Individual Session Control  5938
5         32          Reflect Octets Capability  6038
6         64          Symmetrical Size Sender Test Packet Format  6038
=====

```

ldp-treetrace

Syntax `ldp-treetrace [prefix ip-prefix/mask] [detail]`

Context `show>test-oam`

Description 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.

Output

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
```

```
*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                Failed Hops      : 0
Probe State      : OK                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                Failed Hops      : 0
Probe State      : OK                Failed Probes    : 0
```

```
=====
Discovered Paths
```

```
=====
PathDest      Egr-NextHop      Remote-RtrAddr      Discovery-time
  DiscoveryTtl   ProbeState        ProbeTmOutCnt       RtnCode
-----
127.1.0.5      7      10.10.1.2      12.12.12.10      12/19/2006 05:11:01
                                0      EgressRtr
127.1.0.9      7      10.10.1.2      12.12.12.10      12/19/2006 05:11:01
                                0      EgressRtr
127.1.0.15     7      10.10.1.2      12.12.12.10      12/19/2006 05:11:01
                                0      EgressRtr
127.1.0.19     7      10.10.1.2      12.12.12.10      12/19/2006 05:11:01
                                0      EgressRtr
127.1.0.24     7      10.10.1.2      12.12.12.10      12/19/2006 05:11:01
                                0      EgressRtr
127.1.0.28     7      10.10.1.2      12.12.12.10      12/19/2006 05:11:01
                                0      EgressRtr

.....

127.1.0.252    7      10.10.1.2      12.12.12.10      12/19/2006 05:11:01
                                0      EgressRtr
127.1.0.255    7      10.10.1.2      12.12.12.10      12/19/2006 05:11:01
                                0      EgressRtr
=====
```

```
*A:ALA-48#
```

```
*A:ALA-48# show test-oam twamp server
```

```
=====
TWAMP Server (port 862)
```

```
=====
Admin State : Up                Oper State : Up
Up Time     : 0d 00:00:05
Curr Conn   : 1                Max Conn   : 32
ConnTimeout : 1800             Conn Reject : 2
Curr Sess   : 2                Max Sess   : 32
Tests Done  : 5                Tests Rej   : 0
Tests Abort : 0
TstPktsRx  : 999              TstPktsTx  : 999
=====
```

```
Prefix      : 10.0.0.0/8
Description  : NMS-West
=====
```

```
Admin State : Up                Oper State : Up
Curr Conn    : 1                Max Conn   : 32
Conn Reject  : 0
Curr Sess    : 2                Max Sess   : 32
```

```

Tests Done   : 5
Tests Abort  : 0
TstPktsRx   : 999
Tests Rej    : 0
TstPktsTx   : 999
-----
Client        Sessions      Idle    TstPktsRx  TstPktsTx
              Curr/Done/Rej/Abort
-----
10.1.1.1      2/5/0/0        920     999        999
=====
Prefix        : 10.0.0.0/16
Description    : NMS-West-Special
=====
Admin State   : Up
Curr Conn    : 0
Conn Reject   : 0
Curr Sess    : 0
Tests Done   : 0
Tests Abort  : 0
TstPktsRx    : 0
Oper State   : Up
Max Conn     : 32
Max Sess     : 32
Tests Rej    : 0
TstPktsTx    : 0
-----
Client        Sessions      Idle    TstPktsRx  TstPktsTx
              Curr/Done/Rej/Abort
-----
=====

```

twamp-light

Syntax **twamp-light**

Context show>test-oam>twamp

Description This command enables the context to display TWAMP-Light information.

reflectors

Syntax **reflectors**

Context show>test-oam>twamp>twamp-light

Description This command shows TWAMP-Light reflector information.

Output

Sample Output

```

show test-oam twamp twamp-light reflectors
=====
TWAMP-Light Reflectors
=====
Router/VRN   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-light

Syntax	twamp-light
Context	show>router show>service
Description	This command shows TWAMP-Light reflector information, either for the base router or for a specific service.
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 Reflector Prefixes
-----
Prefix                Description
-----
10.2.1.1/32           Process only 10.2.1.1 TWAMP Light
-----

```

```
172.16.1.0/24                                     Packets
                                                    Process all 172.16.1.0 TWAMP
                                                    Light packets
-----
No. of TWAMP-Light Reflector Prefixes: 2
-----
```

eth-cfm

Syntax	eth-cfm
Context	show
Description	This command enables the context to display CFM information.

association

Syntax	association [<i>ma-index</i>] [detail]
Context	show>eth-cfm
Description	This command displays eth-cfm association information.
Parameters	<i>ma-index</i> — Specifies the MA index. Values 1 to 4294967295 detail — Displays detailed information for the eth-cfm association.

Output

Sample Output

```
ALU-IPD# show eth-cfm association
=====
CFM Association Table
=====
Md-index   Ma-index   Name                               CCM-intrvl Hold-time Bridge-id
-----
3           1          03-0000000100                     1          n/a      100
10          1          FacilityPrt01                     1          n/a      none
=====
ALU-IPD#
```

cfm-stack-table

Syntax	cfm-stack-table cfm-stack-table [{ all-ports all-sdps all-virtuals }] [level 0..7] [direction up down] cfm-stack-table port <i>port-id</i> [vlan <i>qtag</i> [<i>qtag</i>]] [level 0..7] [direction up down]

cfm-stack-table *sdp sdp-id[:vc-id]* [level 0..7] [direction up | down]
cfm-stack-table **virtual** *service-id* [level 0..7]
cfm-stack-table **facility** [{all-ports | all-lags | all-lag-ports | all-tunnel-meps | all-router-interfaces}] [level 0..7] [direction up | down]
cfm-stack-table **facility** **collect-imm-stats**
cfm-stack-table **facility** **lag** *id* [tunnel 1..4094] [level 0..7] [direction up | down]
cfm-stack-table **facility** **port** *id* [level 0..7] [direction up | down]
cfm-stack-table **facility** **router-interface** *ip-int-name* [level 0..7] [direction up | down]

Context show>eth-cfm

Description 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.

level — Display the MD level of the maintenance point.

Values 0 to 7

direction up | down — Displays the direction in which the MP faces on the bridge port.

facility — Displays the CFM stack table information for facility MEPs. The base command will display all the facility MEPs. Options may be included in order to further parse the table for specific facility MEP information.

sdp *sdp-id[:vc-id]* — Displays CFM stack table information for the specified SDP.

virtual *service-id* — Displays CFM stack table information for the specified SDP.

Output

Sample Output

```

show eth-cfm cfm-stack-table
=====
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
A = AisRx, L = CSF LOS Rx, F = CSF AIS/FDI rx, r = CSF RDI rx

=====
CFM SAP Stack Table
=====
Sap                Lvl Dir Md-index  Ma-index  MepId  Mac-address  Defect
-----
1/1/6:20.0         4 B      14      803  MIP  d8:1c:01:01:00:06  -----
1/1/6:3000.1001    4 B      14      800  MIP  00:00:00:00:00:28  -----
1/1/6:2000.1002    4 B      14      802  MIP  d8:1c:01:01:00:06  -----
1/1/6:0.*          4 B      14      805  MIP  d8:1c:01:01:00:06  -----
1/1/9:300          2 U      12      300  28  00:00:00:00:00:28  -----
1/1/9:401          2 U      12      401  28  00:00:00:00:00:28  -----
1/1/9:600          2 U      12      600  28  00:00:00:00:00:28  -----
  
```



```

1/1/9:600          5 B          15          666 MIP 00:10:11:00:00:1c -----
1/1/10:4.*         2 U          12           4  28 00:00:00:00:00:28 --C----
1/1/10:1000.*      5 U          15         1000 28 00:00:00:00:00:28 -----
1/1/10:1001.*      5 U          15         1001 28 00:00:00:00:00:28 -----
=====

```

=====

CFM Ethernet Tunnel Stack Table

=====

Eth-tunnel	Lvl	Dir	Md-index	Ma-index	MepId	Mac-address	Defect

No Matching Entries

=====

=====

CFM Ethernet Ring Stack Table

=====

Eth-ring	Lvl	Dir	Md-index	Ma-index	MepId	Mac-address	Defect

No Matching Entries

=====

=====

CFM Facility Port Stack Table

=====

Port	Tunnel	Lvl	Dir	Md-index	Ma-index	MepId	Mac-address	Defect

1/2/4	0	0	D	10	1	28	00:00:00:00:00:28	-----

=====

=====

CFM Facility LAG Stack Table

=====

Lag	Tunnel	Lvl	Dir	Md-index	Ma-index	MepId	Mac-address	Defect

No Matching Entries

=====

=====

CFM Facility Tunnel Stack Table

=====

Port/Lag	Tunnel	Lvl	Dir	Md-index	Ma-index	MepId	Mac-address	Defect

No Matching Entries

=====

=====

CFM Facility Interface Stack Table

=====

Interface	Lvl	Dir	Md-index	Ma-index	MepId	Mac-address	Defect

v28-v33	1	D	11	1	28	00:00:00:00:00:28	-----

=====

=====

CFM SAP Primary VLAN Stack Table

=====

Sap	Primary VlanId	Lvl	Dir	Md-index	Ma-index	MepId	Mac-address	Defect

```

-----
1/1/6:20.*
      21          4 B          14          804  MIP d8:1c:01:01:00:06  -----
=====

=====
CFM SDP Stack Table
=====
Sdp              Lvl Dir Md-index   Ma-index   MepId  Mac-address      Defect
-----
1:1000           4 D          14         1000      28 00:00:00:00:00:28  -----
2:777            4 D          14         777       28 d8:1c:ff:00:00:00  -----
400:800          4 B          14         800      MIP 00:00:00:00:01:28  -----
=====

=====
CFM Virtual Stack Table
=====
Service          Lvl Dir Md-index   Ma-index   MepId  Mac-address      Defect
-----
No Matching Entries
=====

```

domain

- Syntax** **domain** [*md-index*] [**association** *ma-index* | **all-associations**] [**detail**]
- Context** show>eth-cfm
- Description** This command displays domain information.
- Parameters**
- 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.
 - all-associations** — Displays all associations to the MD.
 - detail** — Displays detailed domain information.

Output

Sample Output

```

*A:node-1# show eth-cfm domain
=====
CFM Domain Table
=====
Md-index   Level Name                               Format
-----
1    4test-1   charString
2           5
25   7    AA:BB:CC:DD:EE:FF-1             macAddressAndUint
=====

```

mep

Syntax	mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [loopback] [linktrace] mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> [remote-mepid <i>mep-id</i>] all-remote-mepids mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> eth-test-results [remote-peer <i>mac-address</i>] mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> one-way-delay-test [remote-peer <i>mac-address</i>] mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> two-way-delay-test [remote-peer <i>mac-address</i>] mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i> two-way-slm-test [remote-peer <i>mac-address</i>]
Context	show>eth-cfm
Description	This command displays Maintenance Endpoint (MEP) information.
Parameters	domain <i>md-index</i> — Displays the index of the MD to which the MP is associated, or 0, if none. association <i>ma-index</i> — 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. remote-mepid — Includes specified remote MEP ID information for the specified MEP. one-way-delay-test — Includes specified MEP information for one-way-delay-test. two-way-delay-test — Includes specified MEP information for two-way-delay-test. two-way-slm-test — Includes specified MEP information for two-way-slm-test. eth-test-results — Include eth-test-result information for the specified MEP. all-remote-mepids — Includes all remote mep-id information for the specified MEP.

Output

Sample Output

```
# show eth-cfm mep 101 domain 3 association 1
=====
Eth-Cfm MEP Configuration Information
=====
Md-index           : 3                      Direction         : Down
Ma-index           : 1                      Admin             : Enabled
MepId              : 101                    CCM-Enable        : Enabled
IfIndex            : 1342177281              PrimaryVid        : 6553700
Description         : (Not Specified)
EngState           : fngReset                ControlMep        : False
LowestDefectPri     : macRemErrXcon          HighestDefect     : none
Defect Flags       : None
Mac Address         : d0:0d:1e:00:01:01      ControlMep        : False
CcmLtmPriority      : 7
```

```

CcmTx           : 19886
Fault Propagation : disabled
MA-CcmInterval  : 1
Eth-IDm Threshold : 3(sec)
Eth-Ais:        : Enabled
Eth-Ais Tx Priorit*: 7
Eth-Ais Tx Interva*: 1
Eth-Ais Tx Levels : 5
Eth-Tst:        : Disabled
CcmSequenceErr  : 0
FacilityFault   : n/a
MA-CcmHoldTime  : 0ms
MD-Level        : 3
Eth-Ais Rx Ais: : No
Eth-Ais Rx Interv*: 1
Eth-Ais Tx Counte*: 388

```

```

Redundancy:
  MC-LAG State : active

```

```

CcmLastFailure Frame:
  None

```

```

XconCcmFailure Frame:
  None

```

```

=====
show eth-cfm mep <mep-id> domain <md-index> association <ma-index> all-remote-
mepids detail

```

```

show eth-cfm mep 28 domain 14 association 2 all-remote-mepids detail

```

```

=====
Eth-CFM Remote-MEP Information
=====

```

```

Remote MEP ID   : 30
Auto Discovered : True
Port Status TLV : Up
MAC Address     : 00:00:00:00:00:30
Chass. ID SubType: local
Chassis ID      : access-012-west
Man Addr Domain : (Not Specified)

State           : True/Grace
RDI             : False
I/F Status TLV : Up
CCM Last Change : 02/06/2014 21:37:00

```

```

Remote MEP ID   : 32
Auto Discovered : True
Port Status TLV : Up
MAC Address     : 00:00:00:00:00:32
Chass. ID SubType: chassisComponent
Chassis ID      : (Not Specified)
Man Addr Domain : (Not Specified)

State           : True/Grace
RDI             : False
I/F Status TLV : Up
CCM Last Change : 02/06/2014 21:37:00

```

```

=====
show eth-cfm mep <mep-id> domain <md-index> association <ma-index> {all-remote-
mepids | remote-mepid <mep-id>} detail

```

```

show eth-cfm mep 28 domain 14 association 2 remote-mepid 30 detail

```

```

=====
Eth-CFM Remote-MEP Information
=====

```

```

Remote MEP ID   : 30
Auto Discovered : True
Port Status TLV : Up
MAC Address     : 00:00:00:00:00:30
Chass. ID SubType: local
Chassis ID      : access-012-west
Man Addr Domain : (Not Specified)

State           : True/Grace
RDI             : False
I/F Status TLV : Up
CCM Last Change : 02/06/2014 21:37:00

```

```
=====
show eth-cfm mep 28 domain 14 association 2 remote-mepid 30
=====
Eth-CFM Remote-Mep Table
=====
R-mepId AD Rx CC RxRdi Port-Tlv If-Tlv Peer Mac Addr      CCM status since
-----
30      T  True  False Up      Up      00:00:00:00:00:30 02/06/2014 21:37:00
=====
Entries marked with a 'T' under the 'AD' column have been auto-discovered.
```

mip

Syntax	mip
Context	show>eth-cfm
Description	This command displays SAPs/bindings provisioned for allowing the default MIP creation.
Output	

Sample Output

```
*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
=====
```

statistics

Syntax	statistics
Context	show>eth-cfm
Description	This command displays the ETH-CFM statistics counters.
Output	

Sample Output

```

show eth-cfm statistics
=====
ETH-CFM System Statistics
=====
Rx Count           : 58300           Tx Count           : 46723
Dropped Congestion : 0               Discarded Error     : 0
=====

Rx Count:                               PPS ETH-CFM CPU Receive Rate
Tx Count:                               PPS ETH-CFM CPU Transmit Rate
Dropped Congestion:                     Valid/Supported ETH-
CFM packets not processed by the CPU as a result of resource contention
Discarded Error:                         Invalid/Malformed/Unsupported ETH-
CFM packets discarded by the CPU

```

system-config

Syntax	system-config
Context	show>eth-cfm
Description	This command shows various ETH-CFM system-level configuration parameters under the config>eth-cfm [{ redundancy slm system }] hierarchies and various system capabilities.
Output	

Sample Output

```

show eth-cfm system-config
=====
CFM System Configuration
=====
Redundancy
  MC-LAG Standby MEP Shutdown: false
  MC-LAG Hold-Timer           : 1 second(s)

Synthetic Loss Measurement
  Inactivity Timer            : 100 second(s)

ETH-CCM Grace-Period
  Transmit Enabled            : true

Sender ID Information
  ChassisID Subtype           : chassisComponent

-----
ETH-CFM System Configuration Limits
-----
Component                               Current Usage      System Limit
-----
Maintenance Domain (MD)

```

```
Maintenance Association (MA)
  Extended MA (up to 400 MEPs)
Maintenance Endpoint (MEP)
  One-second MEP
  Sub-second MEP
Alarm Indication Signal (AIS)
Client Signal Fail (CSF)
Primary Vlan Ingress MP
Primary Vlan Egress MP
LMM Stats Enabled
LBM Concurrent Tests
  Multicast LB Tests
LTM Concurrent Tests
Intermediate Point (MIP)
-----
=====
```

system-info

Syntax	system-info
Context	show>eth-cfm
Description	This command displays system-level ETH-CFM information states.
Output	

Sample Output

```
show eth-cfm system-info
=====
CFM System State Information
=====
ETH-CCM Grace-Period          : Inactive
=====
```

3.12.2.2.1 OAM Performance Monitoring and Binning Show Commands

bin-group

Syntax	bin-group <i>bin-group-number</i> [detail]
Context	show>oam-pm
Description	Show the configuration data for one or all OAM Performance Monitoring bin groups.
Parameters	<i>bin-group-number</i> — Specifies an OAM Performance Monitoring bin group. Values 1 to 255

detail — Keyword to display additional exclusion and event monitoring information for the bin group.

Output

Sample Output

```
show oam-pm bin-group
```

```
-----  
Configured Lower Bounds for Delay 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	1	500	250
		2	500	1000	500
		3	1000	1500	1000
		4	2000	2000	1500
		5	3000	2500	2000
		6	4000	3000	2500
		7	5000	3500	3000
		8	5500	4000	3500
		9	6500	4500	4000
3	Up	0	0	0	0
		1	1	500	250
		2	500	1000	500
		3	1000	1500	1000
		4	2000	2000	1500
		5	3000	2500	2000
		6	4000	3000	2500
		7	5000	3500	3000
		8	5500	4000	3500
		9	6500	4500	4000

```
-----  
* indicates that the corresponding row element may have been truncated.
```

```
show oam-pm bin-group 3 detail
```

```
-----  
Configured Lower Bounds for Delay Tests, in microseconds  
-----
```

Group Description	Admin	Bin	FD (us)	FDR (us)	IFDV (us)
3	Up	0	0	0	0
		1	1	500	250
		2	500	1000	500
		3	1000	1500	1000
		4	2000	2000	1500
		5	3000	2500	2000
		6	4000	3000	2500
		7	5000	3500	3000
		8	5500	4000	3500
		9	6500	4500	4000


```
-----  
  
-----  
Bins Excluded from Average  
-----  
Bin Type      Direction      Bins  
-----  
FD            round-trip      0,9  
-----  
  
-----  
Delay Events Configured  
-----  
Bin Type      Direction      Lowest Bin      Lower Bound (us)      Raise      Clear  
-----  
FD            round-trip      8              5500              100       none  
-----  
  
-----  
Bins Excluded from Delay Event Count  
-----  
Bin Type      Direction      Lowest Excluded Bin      Lower Bound (us)  
-----  
FD            round-trip      9              6500  
-----  
  
-----  
Delay Events Configured  
-----  
Bin Type      Direction      Lowest Bin      Lower Bound (us)      Raise      Clear  
-----  
FD            forward        3              1000              200       none  
-----  
  
-----  
Bins Excluded from Delay Event Count  
-----  
Bin Type      Direction      Lowest Excluded Bin      Lower Bound (us)  
-----  
FD            forward        4              2000  
-----
```

bin-group-using

Syntax `bin-group-using [bin-group bin-group-number]`

Context `show>oam-pm`

Description Show the list of sessions configured against one or all OAM Performance Monitoring bin groups.

Parameters *bin-group-number* — Specifies an OAM Performance Monitoring 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      vpls1000-PM-AL5-1/1/9:1000.1000              Act
-----
3              Up      vpls1000-PM-YL4-1/1/9:1000.1000              Act
-----
=====
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      vpls1000-PM-AL5-1/1/9:1000.1000              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** Show the configuration and status information for an OAM Performance Monitoring session.
- Parameters**
- session-name* — Specifies the session name up to 32 characters in length.
 - all** — Displays all attributes.
 - base** — Specifies the base configuration option for the session.
 - bin-group** — Specifies the associated bin group and its attributes.
 - event-mon** — Configures event monitoring and last TCA.
 - meas-interval** — Configures event monitoring and last TCA.

Output**Sample Output**

```

show oam-pm session "vpls1000-PM-YL4-1/1/9:1000.1000" all
-----
Basic Session Configuration

```

```
-----
Session Name       : vpls1000-PM-YL4-1/1/9:1000.1000
Description        : (Not Specified)
Test Family        : ethernet           Session Type       : proactive
Bin Group          : 3
-----
```

Ethernet Configuration

```
-----
Source MEP         : 30                 Priority           : 7
Source Domain      : 14                 Dest MAC Address   : 00:00:00:00:00:32
Source Assoc'n     : 1000
-----
```

DMM Test Configuration and Status

```
-----
Test ID           : 10001               Admin State        : Up
Oper State        : Up                  Data TLV Size      : 0 octets
On-Demand Duration: Not Applicable      On-Demand Remaining: Not Applicable
Interval          : 1000 ms
Detectable Tx Err : none
-----
```

SLM Test Configuration and Status

```
-----
Test ID           : 10001               Admin State        : Up
Oper State        : Up                  Data TLV Size      : 0 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%
Detectable Tx Err : none
-----
```

5-mins Measurement Interval Configuration

```
-----
Duration          : 5-mins              Intervals Stored    : 32
Boundary Type     : clock-aligned        Clock Offset        : 0 seconds
Accounting Policy : none                 Event Monitoring     : disabled
Delay Event Mon   : disabled             Loss Event Mon      : disabled
-----
```

Configured Lower Bounds for Delay Tests, in microseconds

Group Description	Admin	Bin	FD (us)	FDR (us)	IFDV (us)
3	Up	0	0	0	0
		1	1	500	250
		2	500	1000	500
		3	1000	1500	1000
		4	2000	2000	1500
		5	3000	2500	2000
		6	4000	3000	2500
		7	5000	3500	3000

```

8      5500      4000      3500
9      6500      4500      4000

```

```

-----
Bins Excluded from Average
-----

```

```

Bin Type      Direction      Bins
-----
FD            round-trip      0,9
-----

```

```

-----
Bins Excluded from Delay Event Count
-----

```

```

Bin Type      Direction      Lowest Excluded Bin      Lower Bound (us)
-----
FD            round-trip      9                        6500
-----

```

```

-----
Delay Events for the DMM Test
-----

```

```

Bin Type      Direction      LowerBound(us)      Raise      Clear      Last TCA (UTC)
-----
FD            round-trip      5500      100      none      none
-----

```

```

-----
Loss Events for the SLM Test
-----

```

```

Event Type      Direction      Raise      Clear      Last TCA (UTC)
-----

```

```

show oam-pm session "vp1s1000-PM-YL4-1/1/9:1000.1000" base

```

```

-----
Basic Session Configuration
-----

```

```

Session Name      : vp1s1000-PM-YL4-1/1/9:1000.1000
Description       : (Not Specified)
Test Family       : ethernet      Session Type      : proactive
Bin Group         : 3

```

```

-----
Ethernet Configuration
-----

```

```

Source MEP        : 30      Priority          : 7
Source Domain     : 14      Dest MAC Address  : 00:00:00:00:00:32
Source Assoc'n    : 1000

```

```

-----
DMM Test Configuration and Status
-----

```

```

Test ID          : 10001      Admin State       : Up
Oper State       : Up        Data TLV Size     : 0 octets
On-Demand Duration: Not Applicable  On-Demand Remaining: Not Applicable

```

Interval : 1000 ms
Detectable Tx Err : none

SLM Test Configuration and Status

Test ID : 10001	Admin State : Up
Oper State : Up	Data TLV Size : 0 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%
Detectable Tx Err : none	

show oam-pm session "vpls1000-PM-YL4-1/1/9:1000.1000" bin-group

Configured Lower Bounds for Delay Tests, in microseconds

Group Description	Admin	Bin	FD (us)	FDR (us)	IFDV (us)
3	Up	0	0	0	0
		1	1	500	250
		2	500	1000	500
		3	1000	1500	1000
		4	2000	2000	1500
		5	3000	2500	2000
		6	4000	3000	2500
		7	5000	3500	3000
		8	5500	4000	3500
		9	6500	4500	4000

Bins Excluded from Average

Bin Type	Direction	Bins
FD	round-trip	0,9

Bins Excluded from Delay Event Count

Bin Type	Direction	Lowest Excluded Bin	Lower Bound (us)
FD	round-trip	9	6500

show oam-pm session "vpls1000-PM-YL4-1/1/9:1000.1000" event-mon

Bins Excluded from Delay Event Count

Bin Type	Direction	Lowest Excluded Bin	Lower Bound (us)
FD	round-trip	9	6500

Delay Events for the DMM Test

Bin Type	Direction	LowerBound(us)	Raise	Clear	Last TCA (UTC)
FD	round-trip	5500	100	none	none

Loss Events for the SLM Test

Event Type	Direction	Raise	Clear	Last TCA (UTC)

```
show oam-pm session "vpls1000-PM-YL4-1/1/9:1000.1000" meas-interval
```

5-mins Measurement Interval Configuration

Duration	: 5-mins	Intervals Stored	: 32
Boundary Type	: clock-aligned	Clock Offset	: 0 seconds
Accounting Policy	: none	Event Monitoring	: disabled
Delay Event Mon	: disabled	Loss Event Mon	: disabled

```
show oam-pm statistics session "eth-pm-service-1000" lmm meas-interval 15-  
mins interval-number 2
```

Start (UTC)	: 2014/07/08 03:15:00	Status	: completed
Elapsed (seconds)	: 900	Suspect	: no
Frames Sent	: 90	Frames Received	: 90

Data Frames Sent Data Frames Received

Forward	900	900
Backward	18900	18900

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	90	0	0	0	0	0
Backward	90	0	0	0	0	0

```

                Und-Delta-T
-----
Forward          0
Backward         0
-----

```

sessions

- Syntax** `sessions [test-family {ethernet | ip}] {event-mon | detectable-tx-errors}`
- Context** `show>oam-pm`
- Description** This command shows a summary of the OAM Performance Monitoring sessions.
- Parameters**
- test-family** — Keyword to show all sessions that match the specified test family type when an optional filter is included.
 - ethernet** — Keyword to specify Ethernet session types.
 - ip** — Keyword to specify IP session types.
 - event-mon** — Keyword to provide a summary of all event monitoring and current state for each session.
 - detectable-tx-errors** — Keyword to provide a summary of tests with detectable transmission errors that prevent the test from sending packets. Not all errors are detectable.

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
-----
vppls1000-PM-AL5-1/1/9:1000.1000     Act      2    proactive  DMM        SLM
vppls1000-PM-YL4-1/1/9:1000.1000     Act      3    proactive  DMM        SLM
=====

OAM Performance Monitoring Session Summary for the IP Test Family
=====
Session                               State   Bin Group   Sess Type   Test Types
-----

show oam-pm sessions event-mon
=====
OAM Performance Monitoring Event Summary for the Ethernet Test Family
=====
Event Monitoring Table Legend:
F = Forward,  B = Backward,  R = Round Trip,  A = Aggregate,

```

```

- = Threshold Not Config,  c = Threshold Config,  * = TCA Active,  P = Pending
=====
Session                                Test   FD FDR IFDV FLR CHLI HLI UNAV UDAV UDUN
Type   FBR FBR  FBR  FB  FBA FBA  FBA  FBA  FBA
-----
vpls1000-PM-AL5-1/1/9:1000.1000    DMM   --- ---  ---
vpls1000-PM-AL5-1/1/9:1000.1000    SLM
vpls1000-PM-YL4-1/1/9:1000.1000    DMM   --c ---  ---
vpls1000-PM-YL4-1/1/9:1000.1000    SLM
=====

=====
OAM Performance Monitoring Event Summary for the IP Test Family
=====
Event Monitoring Table Legend:
F = Forward,  B = Backward,  R = Round Trip,  A = Aggregate,
- = Threshold Not Config,  c = Threshold Config,  * = TCA Active,  P = Pending
=====
Session                                Test   FD FDR IFDV FLR CHLI HLI UNAV UDAV UDUN
Type   FBR FBR  FBR  FB  FBA FBA  FBA  FBA  FBA
-----

show oam-pm sessions detectable-tx-errors
=====
OAM Performance Monitoring Transmit Error Summary: Ethernet Test Family
=====
Session                                Test
Type                                   Detectable Transmit Error
-----
vpls1000-PM-YL4-1/1/9:1000.1000    DMM                                   MEP is administratively down
vpls1000-PM-YL4-1/1/9:1000.1000    SLM                                   MEP is administratively down
=====

=====
OAM Performance Monitoring Transmit Error Summary: IP Test Family
=====
Session                                Test
Type                                   Detectable Transmit Error
-----

```

statistics

Syntax	statistics session <i>session-name</i> {dmm lmm slm twamp-light} meas-interval {raw 5-mins 15-mins 1-hour 1-day} [{all bins summary}] interval-number <i>interval-number</i> [{delay loss}]
Context	show>oam-pm
Description	Show OAM Performance Monitoring delay or loss statistics. The twamp-light test type only shows statistics that are being recorded. If both delay and loss statistics are being recorded for a test and the test is shut down, and the record-stats command parameter is then modified, only statistics for the current record-stats parameter configuration can be displayed.

When a **no shutdown** command is entered for any test, the memory is allocated to that test and all stored results for the previous iteration of that test are deleted.

Output

Sample Output

```
show oam-pm statistics session "vpls1000-PM-YL4-1/1/9:1000.1000" dmm meas-  
interval 5-  
mins interval-number 2
```

```
-----  
Start (UTC)      : 2016/01/06 02:40:00      Status      : completed  
Elapsed (seconds) : 300                      Suspect      : no  
Frames Sent      : 300                      Frames Received : 300  
-----
```

```
-----  
Bin Type   Direction   Minimum (us)   Maximum (us)   Average (us)   EfA  
-----  
FD         Forward      0              5160           784           no  
FD         Backward     0              12554          1432          no  
FD         Round Trip   1266           7948           1966          yes  
FDR        Forward      0              5160           776           no  
FDR        Backward     0              12392          1410          no  
FDR        Round Trip   0              6591           621           no  
IFDV       Forward      0              5071           675           no  
IFDV       Backward     0              10323          954           no  
IFDV       Round Trip   0              6205           536           no  
-----
```

EfA = yes: one or more bins configured to be excluded from the Average calc.

Frame Delay (FD) Bin Counts

```
-----  
Bin      Lower Bound   Forward   Backward   Round Trip  
-----  
0         0 us          96        22         0  
1         1 us          33        43         0  
2        500 us         51        55         0  
3       1000 us        105       103        211  
4       2000 us         10        62         69  
5       3000 us          3        11         10  
6       4000 us          1         1          4  
7       5000 us          1         0          3  
8       5500 us          0         1          2  
9       6500 us          0         2          1  
-----
```

Frame Delay Range (FDR) Bin Counts

```
-----  
Bin      Lower Bound   Forward   Backward   Round Trip  
-----  
0         0 us          131       68        181  
1        500 us          50        54         78  
2       1000 us          73        49         20  
3       1500 us          31        53          7  
4       2000 us           9        40          3  
5       2500 us           1        21          2  
-----
```

6	3000 us	1	7	3
7	3500 us	2	4	2
8	4000 us	1	0	2
9	4500 us	1	4	2

Inter-Frame Delay Variation (IFDV) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	109	63	133
1	250 us	41	62	80
2	500 us	75	64	47
3	1000 us	41	52	14
4	1500 us	19	32	10
5	2000 us	6	13	3
6	2500 us	1	5	4
7	3000 us	0	1	3
8	3500 us	4	2	3
9	4000 us	3	5	2

show oam-pm statistics session "vpls1000-PM-YL4-1/1/9:1000.1000" slm meas-interval 5
-mins interval-number 2

Start (UTC) : 2016/01/06 02:40:00 Status : completed
Elapsed (seconds) : 300 Suspect : no
Frames Sent : 3000 Frames Received : 3000

Frames Sent Frames Received

Forward 3000 3000
Backward 3000 3000

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 300 0 0 0 0 0
Backward 300 0 0 0 0 0

show oam-pm statistics session "vpls1000-PM-YL4-1/1/9:1000.1000" dmm meas-
interval raw

Start (UTC) : 2016/01/06 02:08:27 Status : in-progress
Elapsed (seconds) : 2433 Suspect : yes
Frames Sent : 1023 Frames Received : 791

Bin Type	Direction	Minimum (us)	Maximum (us)	Average (us)	EfA
FD	Forward	0	5536	715	no
FD	Backward	0	12554	1463	no
FD	Round Trip	1266	7948	1935	yes
FDR	Forward	0	5536	715	no
FDR	Backward	0	12554	1450	no
FDR	Round Trip	0	6602	617	no
IFDV	Forward	0	5536	662	no
IFDV	Backward	0	10323	863	no
IFDV	Round Trip	0	6205	506	no

EfA = yes: one or more bins configured to be Excluded from the Average calc.

Frame Delay (FD) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	290	60	0
1	1 us	99	93	0
2	500 us	118	131	0
3	1000 us	248	284	570
4	2000 us	26	192	176
5	3000 us	6	23	24
6	4000 us	2	4	8
7	5000 us	1	1	8
8	5500 us	1	1	3
9	6500 us	0	2	2

Frame Delay Range (FDR) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	390	156	447
1	500 us	118	133	247
2	1000 us	152	122	49
3	1500 us	96	162	18
4	2000 us	22	120	8
5	2500 us	4	68	4
6	3000 us	4	18	5
7	3500 us	2	5	7
8	4000 us	1	1	3
9	4500 us	3	7	4

Inter-Frame Delay Variation (IFDV) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	296	193	369
1	250 us	119	149	201
2	500 us	168	180	129
3	1000 us	111	118	35
4	1500 us	63	90	15
5	2000 us	17	31	11
6	2500 us	5	17	13
7	3000 us	2	4	5

```

8          3500 us          4          3          8
9          4000 us          6          6          5
-----

```

```

show oam-pm statistics session "vpls1000-PM-YL4-1/1/9:1000.1000" slm meas-
interval raw
-----

```

```

Start (UTC)      : 2016/01/06 01:45:56      Status      : in-progress
Elapsed (seconds) : 3791                    Suspect      : yes
Frames Sent      : 23808                    Frames Received : 20035
-----

```

```

-----
Frames Sent      Frames Received
-----
Forward          19894          19889
Backward         19890          19890
-----

```

Frame Loss Ratios

```

-----
Minimum      Maximum      Average
-----
Forward      0.000%      50.000%      0.013%
Backward     0.000%      0.000%      0.000%
-----

```

Availability Counters (Und = Undetermined)

```

-----
Available      Und-Avail      Unavailable      Und-Unavail      HLI      CHLI
-----
Forward        3749          1758            28              0          0          0
Backward       3777          1785            0              0          0          0
-----

```

```

show oam-pm statistics session "base-1.1.1.32" twamp-light meas-interval 5-mins
interval-number 2 delay
-----

```

```

Start (UTC)      : 2016/01/06 03:05:00      Status      : completed
Elapsed (seconds) : 300                    Suspect      : no
Frames Sent      : 300                    Frames Received : 300
-----

```

TWAMP-LIGHT DELAY STATISTICS

```

-----
Bin Type      Direction      Minimum (us)      Maximum (us)      Average (us)      EfA
-----
FD            Forward          0                10015             278              no
FD            Backward         0                2699              941              no
FD            Round Trip       476             1857              635              no
FDR           Forward          0                10015             278              no
FDR           Backward         0                2699              941              no
FDR           Round Trip       12              1393              171              no
IFDV          Forward          0                10015             365              no
IFDV          Backward         0                2465              752              no
IFDV          Round Trip       0                1296              123              no
-----

```

```

EfA = yes: one or more bins configured to be Excluded from the Average calc.
-----

```

Frame Delay (FD) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	182	54	0
1	1 us	49	42	7
2	500 us	44	70	284
3	1000 us	24	102	9
4	2000 us	0	32	0
5	3000 us	0	0	0
6	4000 us	0	0	0
7	5000 us	0	0	0
8	5500 us	0	0	0
9	6500 us	1	0	0

Frame Delay Range (FDR) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	231	96	288
1	500 us	44	70	10
2	1000 us	20	58	2
3	1500 us	4	44	0
4	2000 us	0	24	0
5	2500 us	0	8	0
6	3000 us	0	0	0
7	3500 us	0	0	0
8	4000 us	0	0	0
9	4500 us	1	0	0

Inter-Frame Delay Variation (IFDV) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	176	82	259
1	250 us	37	40	28
2	500 us	65	87	11
3	1000 us	17	45	2
4	1500 us	3	32	0
5	2000 us	0	14	0
6	2500 us	0	0	0
7	3000 us	0	0	0
8	3500 us	0	0	0
9	4000 us	2	0	0

show oam-pm statistics session "base-1.1.1.32" twamp-light meas-interval 5-mins
interval-number 2 loss

Start (UTC)	: 2016/01/06 03:05:00	Status	: completed
Elapsed (seconds)	: 300	Suspect	: no
Frames Sent	: 300	Frames Received	: 300

TWAMP-LIGHT LOSS STATISTICS

Frames Sent	Frames Received

```
-----
Forward                300                300
Backward               300                300
-----
```

```
-----
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          30          0          0          0          0          0
Backward         30          0          0          0          0          0
-----
```

```
=====
```

3.12.2.3 Clear Commands

saa

Syntax	saa-test [<i>test-name</i> [owner <i>test-owner</i>]]
Context	clear
Description	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.
Parameters	<i>test-name</i> — Specifies the name of the SAA test. The test name must already be configured in the config>saa>test context. owner <i>test-owner</i> — 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”.

session

Syntax	session <i>session-name</i> { dmm lmm slm twamp-light }
Context	clear>oam-pm
Description	This command clears the raw measurement interval for the specified session and test.

auto-discovered-meps

Syntax	auto-discovered-meps [<i>mep-id</i>] domain <i>md-index</i> association <i>ma-index</i>
Context	clear>eth-cfm
Description	This clear command provides the necessary mechanism to clear a remote MEP that was auto discovered. The function will clear a specific auto-discovered MEP learned within an association or all auto-discovered MEPs in the association. When the <i>mep-id</i> representing the auto-discovered MEP is omitted and only the domain <i>md-index</i> and association <i>ma-index</i> are provided, ALL auto-discovered MEPs in the association will be cleared. At a minimum the domain <i>md-index</i> and the association <i>ma-index</i> must be provided. Only auto-discovered MEPs may be cleared. This command has no affect on manually configured MEPs.
Default	clear all auto discovered MEP IDs

Parameters	<i>mep-id</i> — Specifies the MEP ID of the remote MEP that was auto-discovered. Values 1 to 8191 <i>md-index</i> — Specifies the domain context in which the remote MEP was auto-discovered. Values 1 to 4294967295 <i>ma-index</i> — Specifies the association context in which the remote MEP was auto-discovered. Values 1 to 4294967295

statistics

Syntax	statistics
Context	clear>eth-cfm
Description	This command clears the eth-cfm statistics counters maintained in clearEthCfmStatistics.

3.12.2.4 Monitor Commands

session

Syntax	session <i>session-name</i> { dmm lmm slm twamp-light }
Context	monitor>oam-pm
Description	This command monitors the raw measurement interval for the specified session and test.
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
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 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)
-----
```

```
-----
      Available  Und-Avail  Unavailable  Und-Unavail      HLI      CHLI
-----
Forward          10         0           0           0           0         0
Backward         10         0           0           0           0         0
-----
```

```
monitor oam-pm session "ip-vprn-500" twamp-light
```

```
-----
At time t = 0 sec (Base Statistics)
-----
```


Frame Delay (FD) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	89719	113813	82529
1	1000 us	51728	43288	62811
2	2000 us	19304	7882	16979
3	3000 us	5207	1300	3067
4	4000 us	1166	335	1280
5	5000 us	469	255	781
6	6000 us	227	129	361
7	7000 us	121	166	152
8	8000 us	83	253	114
9	10000 us	125	728	75

Frame Delay Range (FDR) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	167124	166618	167138
1	5000 us	1025	1531	1011

Inter-Frame Delay Variation (IFDV) Bin Counts

Bin	Lower Bound	Forward	Backward	Round Trip
0	0 us	29284	45291	36062
1	100 us	9615	10793	28238
2	200 us	9289	9827	20379
3	300 us	8933	8733	14325
4	400 us	8597	8362	10257
5	500 us	8216	7789	7635
6	600 us	8178	7606	5893
7	700 us	7782	7345	4963
8	800 us	14799	14500	8416
9	1000 us	63455	47902	31980

3.12.2.5 Debug Commands

eth-cfm

Syntax	eth-cfm
Context	debug
Description	This command enables the context to configure ETH-CFM debugging functions.

mep

Syntax	[no] mep <i>mep-id</i> domain <i>md-index</i> association <i>ma-index</i>
Context	debug>eth-cfm
Description	This command specifies the MEP from which to debug the CFM PDUs. The no form of the command removes the MEP parameters.
Parameters	<i>mep-id</i> — the maintenance association endpoint identifier of the launch point Values 1 to 8191 <i>md-index</i> — the maintenance domain (MD) index value of the launch point Values 1 to 4294967295 <i>ma-index</i> — the maintenance association (MA) index value of the launch point Values 1 to 4294967295

mip

Syntax	[no] mip domain <i>md-index</i> association <i>ma-index</i>
Context	debug>eth-cfm
Description	This command specifies the MIP from which to debug the CFM PDUs. The no form of the command removes the MIP parameters.
Parameters	<i>md-index</i> — the maintenance domain (MD) index value of the launch point Values 1 to 4294967295 <i>ma-index</i> — the maintenance association (MA) index value of the launch point Values 1 to 4294967295

packet

Syntax	packet all packet cfm-opcode <i>opcode</i> [<i>opcode</i> ... (up to 5 max.)] no packet
Context	debug>eth-cfm>mep debug>eth-cfm>mip
Description	This command defines the ETH-CFM opcodes of interest to be debugged. The no form of the command stops packet debugging and the collection of PDUs.
Parameters	all — debugging is enabled for all ETH-CFM packets <i>opcode</i> — specifies the standard numerical reference or common three-letter acronym (TLA) that identifies the CFM PDU type. Up to five opcodes can be specified, and a combination of both numbers and TLAs can be used. MEPs support all opcodes. MIPs support 2 (LBR), 3 (LBM), 4 (LTR), and 5 (LTM). Unknown or unsupported opcodes in TLA form are rejected. The applicable numerical opcode can be used instead. When numerical references are used, they are converted to a known TLA or left in numerical form if the TLA is unknown. Re-entering the packet command will overwrite the previous opcode entries for the MEP or MIP. Values MEP: 1 to 255 common TLA MIP: 2 to 5 common TLA

lsp-ping-trace

Syntax	lsp-ping-trace [tx rx both] [raw detail] no lsp-ping-trace
Context	debug>oam
Description	This command enables debugging for lsp-ping.
Parameters	tx rx both — Specifies to enable LSP ping debugging for TX, RX, or both RX and TX for the for debug direction. raw detail — Displays output for the for debug mode.

3.12.2.6 Tools Commands

debug-packet

Syntax	debug-packet [clear]
Context	tools>dump>eth-cfm
Description	This command displays and optionally clears the counters representing the number of CFM PDUs that matched the debug criteria but were not passed to the debug logger. This situation is caused by a full message queue.
Parameters	clear — clears the current counters.
Output	

Sample Output

```
tools dump eth-cfm debug-packet
=====
ETH-CFM Debug Logging Message Queue Statistics
=====
Rx Debug Exceptions      : 0
Tx Debug Exceptions      : 0
=====
```

top-active-mep

Syntax	top-active-mep [rx-sort tx-sort] [clear]
Context	tools>dump>eth-cfm
Description	This command displays and optionally clears the most active MEPs on the system.
Default	Sorts total in both directions
Parameters	rx-sort — Sort in the RX direction. tx-sort — Sort in the TX direction. clear — clears the current counters.

4 Standards and Protocol Support



Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attr-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-04, *Advertisement of Multiple Paths in BGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*
draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*
draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*
draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*
RFC 1772, *Application of the Border Gateway Protocol in the Internet*
RFC 1997, *BGP Communities Attribute*
RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*
RFC 2439, *BGP Route Flap Damping*
RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*
RFC 2858, *Multiprotocol Extensions for BGP-4*
RFC 2918, *Route Refresh Capability for BGP-4*
RFC 3107, *Carrying Label Information in BGP-4*
RFC 3392, *Capabilities Advertisement with BGP-4*
RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*
RFC 4360, *BGP Extended Communities Attribute*
RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*
RFC 4486, *Subcodes for BGP Cease Notification Message*
RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*
RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/ MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*
RFC 4724, *Graceful Restart Mechanism for BGP (helper mode)*
RFC 4760, *Multiprotocol Extensions for BGP-4*
RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)*
RFC 4893, *BGP Support for Four-octet AS Number Space*
RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*
RFC 5065, *Autonomous System Confederations for BGP*
RFC 5291, *Outbound Route Filtering Capability for BGP-4*
RFC 5575, *Dissemination of Flow Specification Rules*
RFC 5668, *4-Octet AS Specific BGP Extended Community*
RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*
RFC 6811, *Prefix Origin Validation*
RFC 7607, *Codification of AS 0 Processing*

Circuit Emulation

MEF-8, *Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks*, October 2004

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

IEEE 802.1ak, *Multiple Registration Protocol*

IEEE 802.1aq, *Shortest Path Bridging*

IEEE 802.1ax, *Link Aggregation*

IEEE 802.1D, *MAC Bridges*

IEEE 802.1p, *Traffic Class Expediting*

IEEE 802.1Q, *Virtual LANs*

IEEE 802.1s, *Multiple Spanning Trees*

IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*

IEEE 802.1X, *Port Based Network Access Control*

IEEE 802.3ab, *1000BASE-T*

IEEE 802.3ac, *VLAN Tag*

IEEE 802.3ad, *Link Aggregation*

IEEE 802.3ae, *10 Gb/s Ethernet*

IEEE 802.3ah, *Ethernet in the First Mile*

IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*

IEEE 802.3i, *Ethernet*

IEEE 802.3u, *Fast Ethernet*

IEEE 802.3x, *Ethernet Flow Control*

IEEE 802.3z, *Gigabit Ethernet*

ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*

ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*

ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-overlay-02, *A Network Virtualization Overlay Solution using EVPN*

draft-ietf-bess-evpn-prefix-advertisement-02, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-proxy-arp-nd-00, *Operational Aspects of Proxy-ARP/ND in EVPN Networks*

draft-ietf-bess-evpn-vpls-seamless-integ-00, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*

draft-ietf-bess-evpn-vpws-06, *VPWS support in EVPN*

RFC 7432, *BGP MPLS-Based Ethernet VPN*

RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*

Frame Relay

ANSI T1.617 Annex D, *DSS1 - Signalling Specification For Frame Relay Bearer Service*

FRF.1.2, *PVC User-to-Network Interface (UNI) Implementation Agreement*

FRF.12, *Frame Relay Fragmentation Implementation Agreement*

FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*

FRF.5, *Frame Relay/ATM PVC Network Interworking Implementation*

FRF2.2, *PVC Network-to-Network Interface (NNI) Implementation Agreement*

ITU-T Q.933 Annex A, *Additional procedures for Permanent Virtual Connection (PVC) status management*

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 4204, *Link Management Protocol (LMP)*

RFC 4208, *Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model*

RFC 4872, *RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery*

Intermediate System to Intermediate System (IS-IS)

draft-ginsberg-isis-mi-bis-01, *IS-IS Multi-Instance* (single topology)
draft-ietf-isis-mi-02, *IS-IS Multi-Instance*
draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*
draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*
ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*
RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*
RFC 2973, *IS-IS Mesh Groups*
RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*
RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*
RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*
RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information*
RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*
RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*
RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*
RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*
RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*
RFC 5304, *IS-IS Cryptographic Authentication*
RFC 5305, *IS-IS Extensions for Traffic Engineering TE*
RFC 5306, *Restart Signaling for IS-IS (helper mode)*
RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
RFC 5308, *Routing IPv6 with IS-IS*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5310, *IS-IS Generic Cryptographic Authentication*
RFC 6213, *IS-IS BFD-Enabled TLV*
RFC 6232, *Purge Originator Identification TLV for IS-IS*
RFC 6233, *IS-IS Registry Extension for Purges*
RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*

RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*

RFC 7431, *Multicast-Only Fast Reroute*

RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*

draft-ietf-vrrp-unified-spec-02, *Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6*

RFC 768, *User Datagram Protocol*

RFC 793, *Transmission Control Protocol*

RFC 854, *Telnet Protocol Specifications*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1350, *The TFTP Protocol (revision 2)*

RFC 1534, *Interoperation between DHCP and BOOTP*

RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*

RFC 2131, *Dynamic Host Configuration Protocol*

RFC 2132, *DHCP Options and BOOTP Vendor Extensions*

RFC 2347, *TFTP Option Extension*

RFC 2348, *TFTP Blocksize Option*

RFC 2349, *TFTP Timeout Interval and Transfer Size Options*

RFC 2428, *FTP Extensions for IPv6 and NATs*

RFC 2784, *Generic Routing Encapsulation (GRE)*

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 3046, *DHCP Relay Agent Information Option (Option 82)*

RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3596, *DNS Extensions to Support IP version 6*

RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*

RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*

RFC 4251, *The Secure Shell (SSH) Protocol Architecture*
RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)
RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*
RFC 4254, *The Secure Shell (SSH) Connection Protocol*
RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*
RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*
RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (ECDSA)*
RFC 5880, *Bidirectional Forwarding Detection (BFD)*
RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*
RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*
RFC 6398, *IP Router Alert Considerations and Usage (MLD)*
RFC 6528, *Defending against Sequence Number Attacks*
RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*
RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast*
draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*
draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*
draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*
RFC 1112, *Host Extensions for IP Multicasting*
RFC 2236, *Internet Group Management Protocol, Version 2*
RFC 2375, *IPv6 Multicast Address Assignments*
RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*
RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*
RFC 3376, *Internet Group Management Protocol, Version 3*
RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*
RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*
RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

-
- RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*
- RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*
- RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised) (auto-RP groups)*
- RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*
- RFC 4601, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*
- RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*
- RFC 4607, *Source-Specific Multicast for IP*
- RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*
- RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*
- RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*
- RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*
- RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*
- RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*
- RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*
- RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*
- RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*
- RFC 6513, *Multicast in MPLS/BGP IP VPNs*
- RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*
- RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*
- RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*
- RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*
- RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*
- RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*
- RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

Internet Protocol (IP) — Version 4

RFC 791, *Internet Protocol*
RFC 792, *Internet Control Message Protocol*
RFC 826, *An Ethernet Address Resolution Protocol*
RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*
RFC 1812, *Requirements for IPv4 Routers*
RFC 1918, *Address Allocation for Private Internets*
RFC 1981, *Path MTU Discovery for IP version 6*
RFC 2003, *IP Encapsulation within IP*
RFC 2401, *Security Architecture for Internet Protocol*
RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

Internet Protocol (IP) — Version 6

RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*
RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*
RFC 2473, *Generic Packet Tunneling in IPv6 Specification*
RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*
RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*
RFC 3587, *IPv6 Global Unicast Address Format*
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*
RFC 3971, *SEcure Neighbor Discovery (SEND)*
RFC 3972, *Cryptographically Generated Addresses (CGA)*
RFC 4007, *IPv6 Scoped Address Architecture*
RFC 4193, *Unique Local IPv6 Unicast Addresses*
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
RFC 4862, *IPv6 Stateless Address Autoconfiguration (router functions)*
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*
RFC 5007, *DHCPv6 Leasequery*

RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*
RFC 5952, *A Recommendation for IPv6 Address Text Representation*
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*
draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*
RFC 2401, *Security Architecture for the Internet Protocol*
RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*
RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*
RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*
RFC 2406, *IP Encapsulating Security Payload (ESP)*
RFC 2409, *The Internet Key Exchange (IKE)*
RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*
RFC 2560, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*
RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*
RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*
RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*
RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*
RFC 3947, *Negotiation of NAT-Traversal in the IKE*
RFC 3948, *UDP Encapsulation of IPsec ESP Packets*
RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*
RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*
RFC 4301, *Security Architecture for the Internet Protocol*
RFC 4303, *IP Encapsulating Security Payload*
RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*
RFC 4308, *Cryptographic Suites for IPsec*
RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*

RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec*
RFC 4891, *Using IPsec to Secure IPv6-in-IPv4 Tunnels*
RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*
RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*
RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*
RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*
RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*
RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

Label Distribution Protocol (LDP)

draft-ietf-mpls-ldp-ip-pw-capability-09, *Controlling State Advertisements Of Non-negotiated LDP Applications*
draft-ietf-mpls-ldp-ipv6-15, *Updates to LDP for IPv6*
draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*
draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*
draft-pdutta-mpls-mldp-up-redundancy-00, *Upstream LSR Redundancy for Multipoint LDP Tunnels*
draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*
draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*
RFC 3037, *LDP Applicability*
RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol (helper mode)*
RFC 5036, *LDP Specification*
RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*
RFC 5443, *LDP IGP Synchronization*
RFC 5561, *LDP Capabilities*
RFC 5919, *Signaling LDP Label Advertisement Completion*
RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*
RFC 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*
RFC 2661, *Layer Two Tunneling Protocol "L2TP"*
RFC 2809, *Implementation of L2TP Compulsory Tunneling via RADIUS*
RFC 3438, *Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update*
RFC 3931, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*
RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*
RFC 4719, *Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*
RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*

Management

draft-ietf-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*
draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*
draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*
draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*
draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*
draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*
ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*
ianagmplstc-mib, *IANA-GMPLS-TC-MIB*
ianaiftype-mib, *IANAifType-MIB*
ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*
IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*
IEEE8021-PAE-MIB, *IEEE 802.1X MIB*
IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*
LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*
RFC 1157, *A Simple Network Management Protocol (SNMP)*
RFC 1212, *Concise MIB Definitions*

RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*

RFC 1215, *A Convention for Defining Traps for use with the SNMP*

RFC 1724, *RIP Version 2 MIB Extension*

RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*

RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*

RFC 2206, *RSVP Management Information Base using SMIv2*

RFC 2213, *Integrated Services Management Information Base using SMIv2*

RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*

RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*

RFC 2515, *Definitions of Managed Objects for ATM Management*

RFC 2570, *SNMP Version 3 Framework*

RFC 2571, *An Architecture for Describing SNMP Management Frameworks*

RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*

RFC 2573, *SNMP Applications*

RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*

RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*

RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

RFC 2579, *Textual Conventions for SMIv2*

RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

-
- RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*
- RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
- RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4)*
- RFC 3419, *Textual Conventions for Transport Addresses*
- RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*
- RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*
- RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*
- RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*
- RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*
- RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*
- RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*
- RFC 3877, *Alarm Management Information Base (MIB)*
- RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*
- RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*
- RFC 4001, *Textual Conventions for Internet Network Addresses*
- RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*
- RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*
- RFC 4220, *Traffic Engineering Link Management Information Base*
- RFC 4273, *Definitions of Managed Objects for BGP-4*
- RFC 4292, *IP Forwarding Table MIB*
- RFC 4293, *Management Information Base for the Internet Protocol (IP)*
- RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*
- RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*
- RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms (TLS)*
- RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*
- RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*

RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*
RFC 5102, *Information Model for IP Flow Information Export*
RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2* (TLS client, RSA public key)
RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)* (server, unauthenticated mode)
RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*
RFC 6020, *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*
RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*
RFC 6241, *Network Configuration Protocol (NETCONF)*
RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*
RFC 6243, *With-defaults Capability for NETCONF*
RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*
RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*
RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching - Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*
RFC 5921, *A Framework for MPLS in Transport Networks*
RFC 5960, *MPLS Transport Profile Data Plane Architecture*
RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*
RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*
RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*
RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*
RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*
RFC 6478, *Pseudowire Status for Static Pseudowires*
RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

Open Shortest Path First (OSPF)

draft-ietf-ospf-prefix-link-attr-06, *OSPFv2 Prefix/Link Attribute Advertisement*

draft-ietf-ospf-segment-routing-extensions-04, *OSPF Extensions for Segment Routing*

RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*

RFC 1765, *OSPF Database Overflow*

RFC 2328, *OSPF Version 2*

RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*

RFC 3509, *Alternative Implementations of OSPF Area Border Routers*

RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart (helper mode)*
RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*
RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*
RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*
RFC 4552, *Authentication/Confidentiality for OSPFv3*
RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*
RFC 4970, *Extensions to OSPF for Advertising Optional Router Capabilities*
RFC 5185, *OSPF Multi-Area Adjacency*
RFC 5187, *OSPFv3 Graceful Restart (helper mode)*
RFC 5243, *OSPF Database Exchange Summary List Optimization*
RFC 5250, *The OSPF Opaque LSA Option*
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*
RFC 5340, *OSPF for IPv6*
RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*
RFC 5838, *Support of Address Families in OSPFv3*
RFC 6987, *OSPF Stub Router Advertisement*

OpenFlow

ONF *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*
draft-ietf-pce-segment-routing-05, *PCEP Extensions for Segment Routing*
draft-ietf-pce-stateful-pce-11, *PCEP Extensions for Stateful PCE*
RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*
RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*
RFC 1661, *The Point-to-Point Protocol (PPP)*
RFC 1662, *PPP in HDLC-like Framing*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*
RFC 1989, *PPP Link Quality Monitoring*
RFC 1990, *The PPP Multilink Protocol (MP)*
RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*
RFC 2153, *PPP Vendor Extensions*
RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*
RFC 2615, *PPP over SONET/SDH*
RFC 2686, *The Multi-Class Extension to Multi-Link PPP*
RFC 2878, *PPP Bridging Control Protocol (BCP)*
RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*
RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points (Gx support as it applies to wireline environment (BNG))*
RFC 3588, *Diameter Base Protocol*
RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*
MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*
MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*
MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*
MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*
RFC 3916, *Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3)*
RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*
RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*
RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*
RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*
RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*
RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*

- RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*
- RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*
- RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*
- RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*
- RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*
- RFC 6073, *Segmented Pseudowire*
- RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*
- RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*
- RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*
- RFC 6718, *Pseudowire Redundancy*
- RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*
- RFC 6870, *Pseudowire Preferential Forwarding Status bit*
- RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*
- RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

- RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*
- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2598, *An Expedited Forwarding PHB*
- RFC 3140, *Per Hop Behavior Identification Codes*
- RFC 3260, *New Terminology and Clarifications for Diffserv*

Resource Reservation Protocol - Traffic Engineering (RSVP-TE)

- draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*
- RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*
RFC 2961, *RSVP Refresh Overhead Reduction Extensions*
RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*
RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*
RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling
Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions
(IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful
restart helper procedures)*
RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic
Engineering (RSVP-TE)*
RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic
Engineering*
RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic
Engineering Tunnels*
RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*
RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic
Engineering*
RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware
MPLS Traffic Engineering*
RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS
Traffic Engineering*
RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*
RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering
(RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*
RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*
RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource
Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*
RFC 5712, *MPLS Traffic Engineering Soft Preemption*
RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering
Networks*

Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*
RFC 2080, *RIPng for IPv6*
RFC 2082, *RIP-2 MD5 Authentication*
RFC 2453, *RIP Version 2*

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ITU-G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*

Timing

GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005*

GR-253-CORE, *SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000*

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions, issued 09/2008*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003*

ITU-T G.8261, *Timing and synchronization aspects in packet networks, issued 04/2008*

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007*

ITU-T G.8264, *Distribution of timing information through packet networks, issued 10/2008*

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization, issued 10/2010*

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014*

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications (estimating the interarrival jitter)*

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses (S2a roaming based on GPRS)*

Customer Document and Product Support



Customer Documentation

[Customer Documentation Welcome Page](#)



Technical Support

[Product Support Portal](#)



Documentation Feedback

[Customer Documentation Feedback](#)

