7750 SR OS OAM and Diagnostics Guide

Software Version: 7750 SR OS 9.0 r1
March 2011
Document Part Number: 93-0181-05-01

Table of Contents

List of Tables

List of Figures

Mirror Services

OAM and SAA

Common CLI Command Descriptions

Preface

About This Guide

This guide describes service mirroring and Operations, Administration and Management (OAM) and diagnostic tools provided by the 7750 SR OS 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.

Audience

This manual is intended for network administrators who are responsible for configuring the 7750 SR-Series routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols, standards, and services described in this manual include the following:

- CLI concepts
- Subscriber services
- Service mirroring
- Operation, Administration and Maintenance (OAM) operations

List of Technical Publications

The 7750 SR documentation set is composed of the following books:

- 7750 SR OS Basic System Configuration Guide

  This guide describes basic system configurations and operations.
- 7750 SROS System Management Guide

  This guide describes system security and access configurations as well as event logging and accounting logs.
- 7750 SROS Interface Configuration Guide

  This guide describes card, Media Dependent Adapter (MDA), and port provisioning.
- 7750 SROS Router Configuration Guide

  This guide describes logical IP routing interfaces and associated attributes such as an IP address, port, link aggregation group (LAG) as well as IP and MAC-based filtering.
- 7750 SROS OS Routing Protocols Guide

  This guide provides an overview of routing concepts and provides configuration examples for RIP, OSPF, IS-IS, BGP, and route policies.
- 7750 SR OS MPLS Guide

  This guide describes how to configure Multiprotocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- 7750 SROS Services Guide

  This guide describes how to configure service parameters such as service distribution points (SDPs), customer information, and user services.
- 7750 SR OS OAM and Diagnostic Guide

  This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- 7750 SR OS Triple Play Guide

  This guide describes Triple Play services and support provided by the 7750 SR7450 ESS7710 SR and presents examples to configure and implement various protocols and services.
- 7750 SROS Quality of Service Guide

  This guide describes how to configure Quality of Service (QoS) policy management.
- OS Multi-Service ISA Guide

  This guide describes services provided by integrated service adapters such as Application Assurance, IPSec, ad insertion (ADI) and Network Address Translation (NAT).

Technical Support

If you purchased a service agreement for your 7750 SR router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, contact your welcome center.

Web: http://www1.alcatel-lucent.com/comps/pages/carrier_support.jhtml

Getting Started

In This Chapter

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

Alcatel-Lucent 7750 SR-Series Services Configuration Process

Table 1 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 1: Configuration Process

Area	Task	Chapter
Diagnostics/ Service verification	Mirroring	Mirror Services on page 17
	Lawful Intercept	Lawful Intercept on page 31
	OAM	OAM and SAA on page 125
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support on page 381

Mirror Services

In This Chapter

This chapter provides information to configure mirroring.

Topics in this chapter include:

Service Mirroring

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

Alcatel-Lucent's service mirroring extends and integrates these capabilities into the network and provides significant operational benefits. Each 7750 SR-Series 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.

Alcatel-Lucent's 7750 SR-Series routers support service-based mirroring. While some Layer 3 switches and routers can mirror on a per-port basis within the device, Alcatel-Lucent 7750 SR-Series 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.

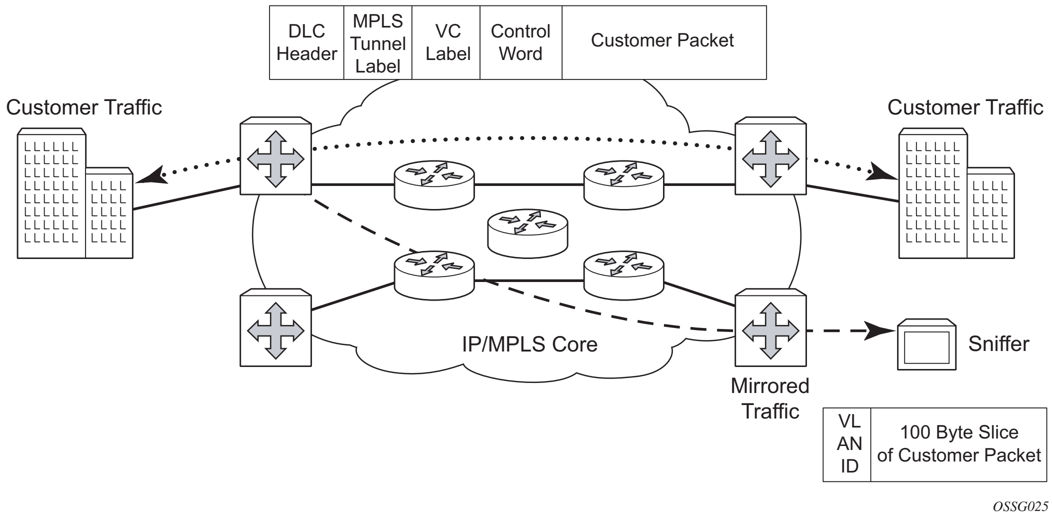

Figure 1: Service Mirroring

Mirror Implementation

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

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

- Ingress and egress packets are mirrored as they appear on the wire. This is important for troubleshooting encapsulation and protocol issues.
  - → When mirroring at ingress, 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 (as seen on the wire) is forwarded to the mirror destination.
- Mirroring must support 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.

Mirror Source and Destinations

Mirror sources and destinations have the following characteristics:

- They can be on the same 7750 SR-Series router (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.

Local and Remote Mirroring

Mirrored frames can be copied and sent to a specific local destination or service on the 7750 SR router (local mirroring) or copies can be encapsulated and sent to a different 7750 SR 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 7750 SR 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 7750 SR router to another through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end 7750 SR 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.

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 7750 SR-Series 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.

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.

Mirroring Performance

Replication of mirrored packets can, typically, affect performance and should be used carefully. Alcatel-Lucent 7750 SR-Series 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. The mirroring architecture also supports mirror rate limiting both at the ingress and egress Flexible Fast Path NPA. This rate limiting is accomplished though a shaping queue and is set according to the maximum amount of mirroring desired.

Mirroring can be performed based on the following criteria:

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

Mirroring Configuration

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

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

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

- Port 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.

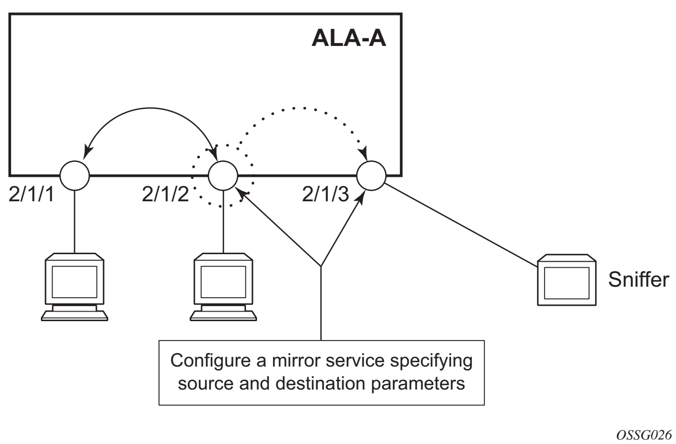

Figure 2: Local Mirroring Example

Figure 3 depicts 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.

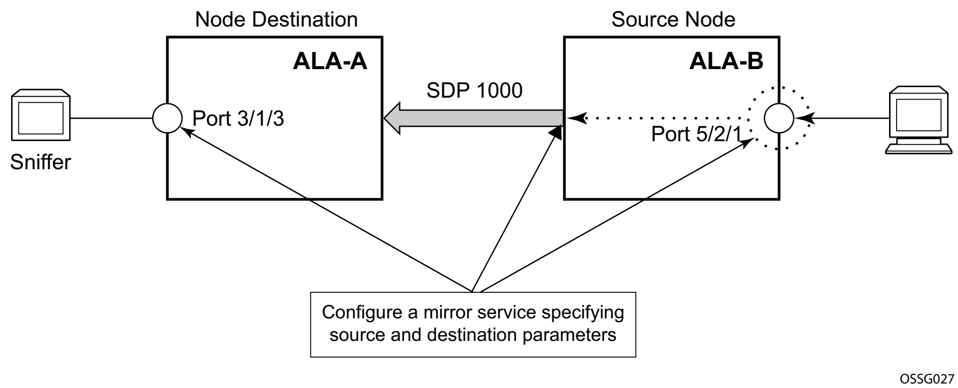

Figure 3: Remote Mirroring Example

ATM Mirroring

ATM mirror functionality allows 7750 SR-Series 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), Ipipe, 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
- Ipipe
- 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.

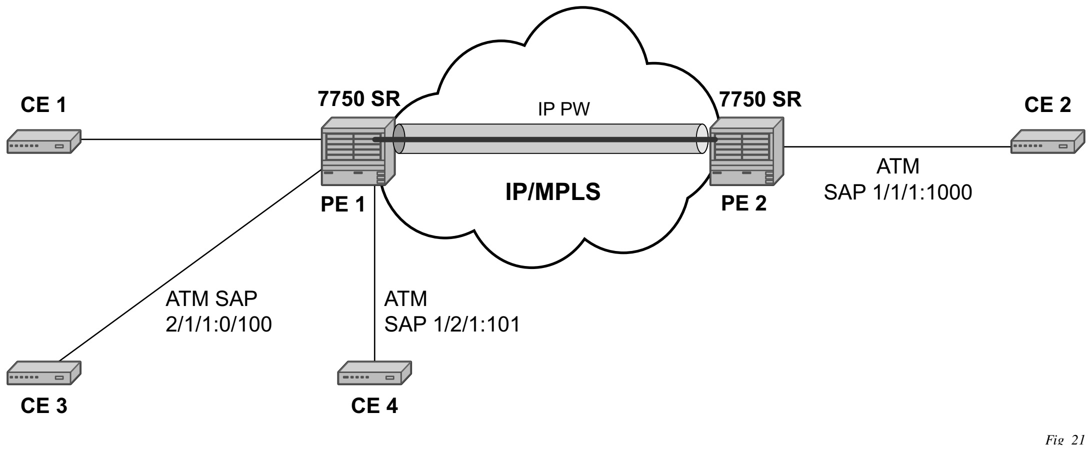

Figure 4: Example of an ATM Mirror Service

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 send towards ATM SAP 1/2/1:1/101.

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.

IP Mirroring

The IP mirroring capability 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, Ipipe services, and subscriber mirrors. It is not supported on VLL services such as Apipe, Epipe, Fpipe, and on ports.

Remote IP Mirroring

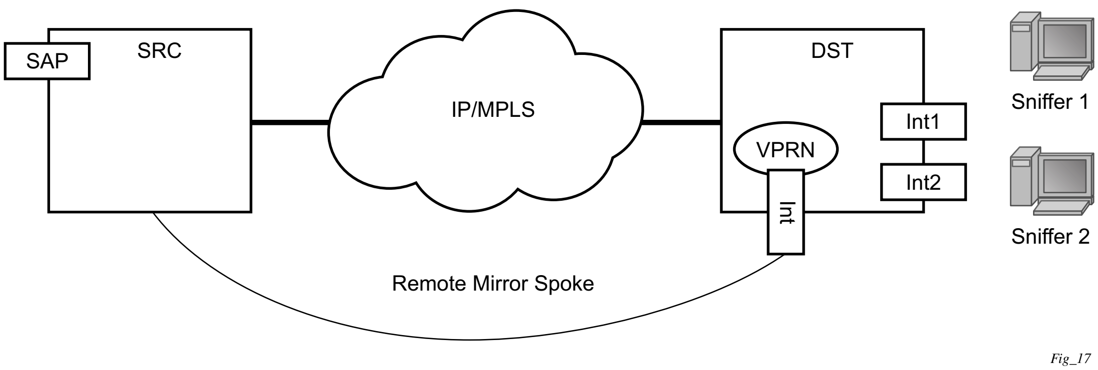

Figure 5: Remote IP Mirroring

With remote IP mirroring, the mirror destination configuration can allow IP packets to be mirrored from a source router (Figure 4). 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.

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.

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.

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.

Subscriber Mirroring

This section describes mirroring based on a subscriber match. 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).

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. Alcatel-Lucent's implementation satisfies most national standard's requirements. LI capability is configurable for all Alcatel-Lucent 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.

LI Activation Via RADIUS

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

Notes: The term "activation" in this section represents both "activation and de-activation".

The activation of an LI session via RADIUS can occur in one of two ways:

- At the time the RADIUS access-accept message is received by the 7x50. In this case, the target (i.e. 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).
- Via RADIUS COA messages. In this case, the target (set of hosts) is identified by either the acct-session-id (which can represent a single host or a collection of hosts) or by a **<sap-id;ip-addr>** carried in the NAS-Port-Id (attr 87) and the Framed-Ip-Address (attr 8).

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 then 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.

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 service** <*service-id*> **active-subscribers li** command shows all sub-hosts with activated LI information. This command will be accessible to cli-user with LI privileges only.

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.

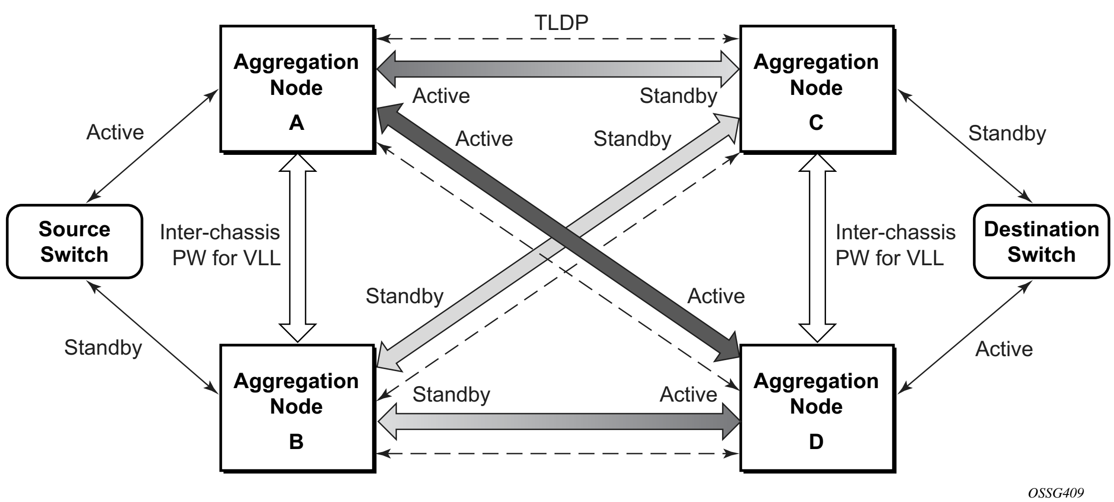

Figure 6: State Engine for Redundant Service to a Redundant Mirror Service

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.

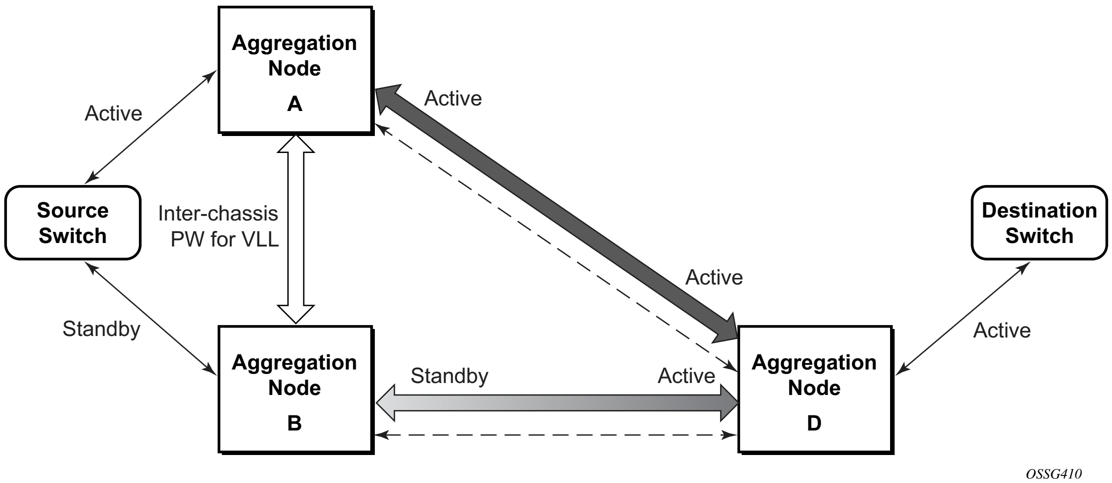

Figure 7: State Engine for Redundant Service to a Non-Redundant Mirror Service

The notable difference with scenarios standard pseudo wire redundancy scenarios is that provided the customer service is redundant on nodes A and B (Figure 5 and Figure 6) 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 5, 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.

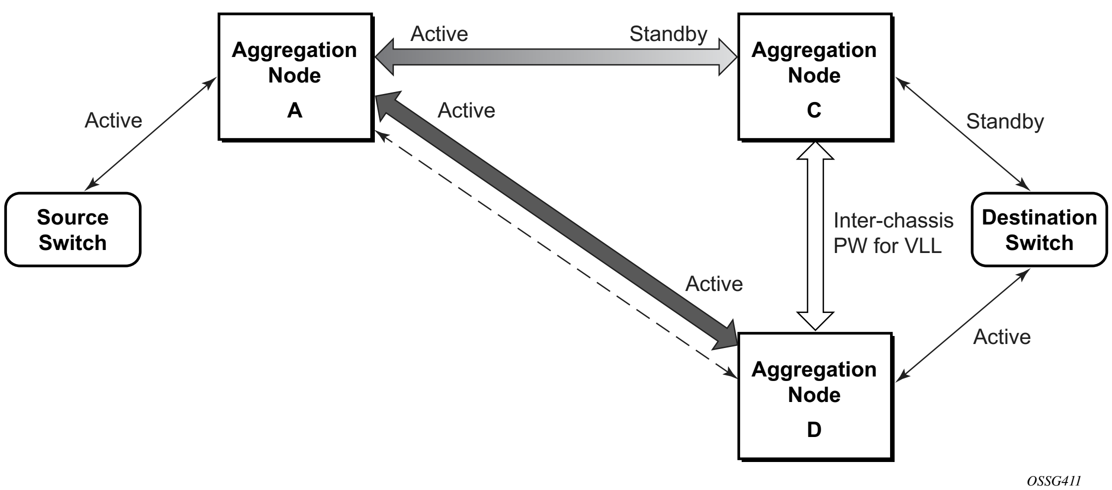

Figure 8: State Engine for a Non-Redundant Service to a Redundant Mirror Service

In the case where a non redundant service is being mirrored to a redundant mirror service (Figure 7) 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).

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.

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.

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:

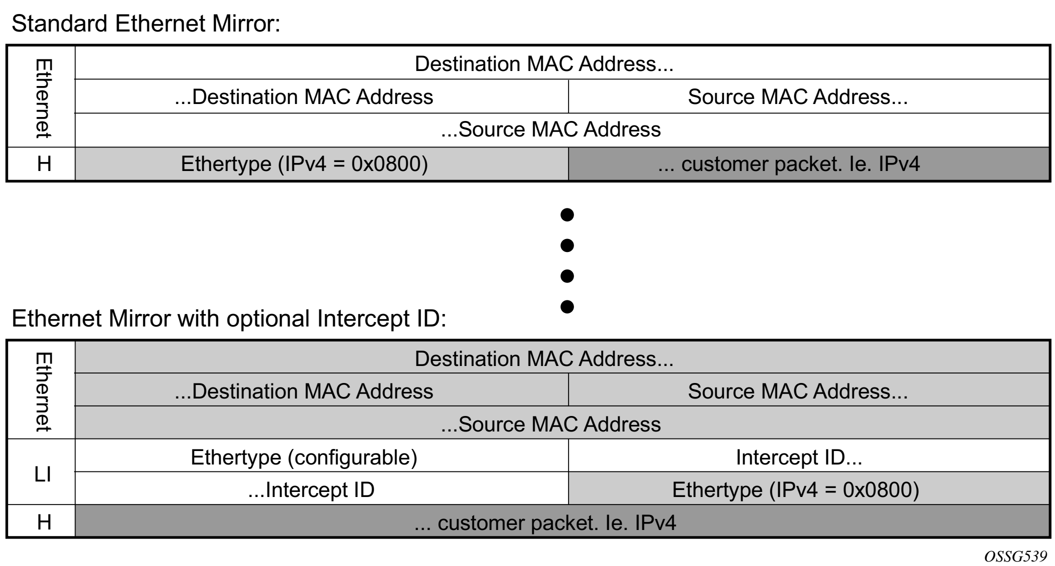

Figure 9: Ethernet Mirror Examples

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]
             ethernet-header [etype hex] [sa mac] [da mac]
```

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.

Configuration Process Overview

Figure 9 displays the process to provision basic mirroring parameters.

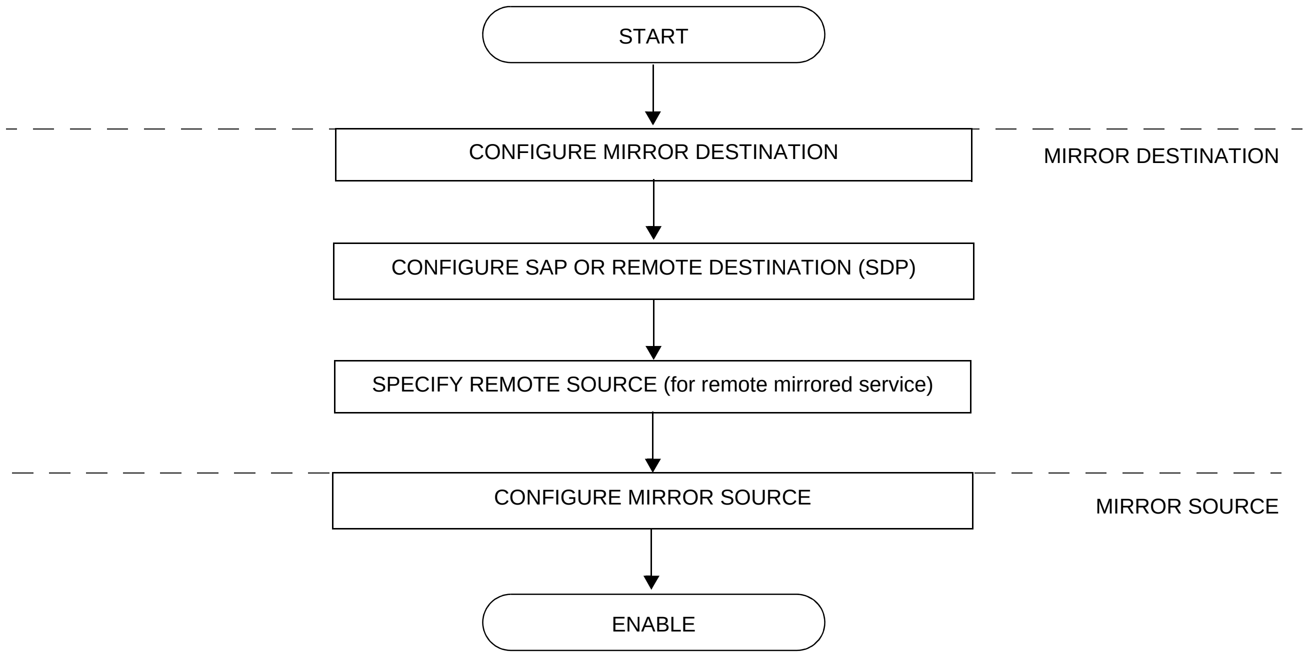

Figure 10: Mirror Configuration and Implementation Flow

Figure 10 displays the process to provision LI parameters.

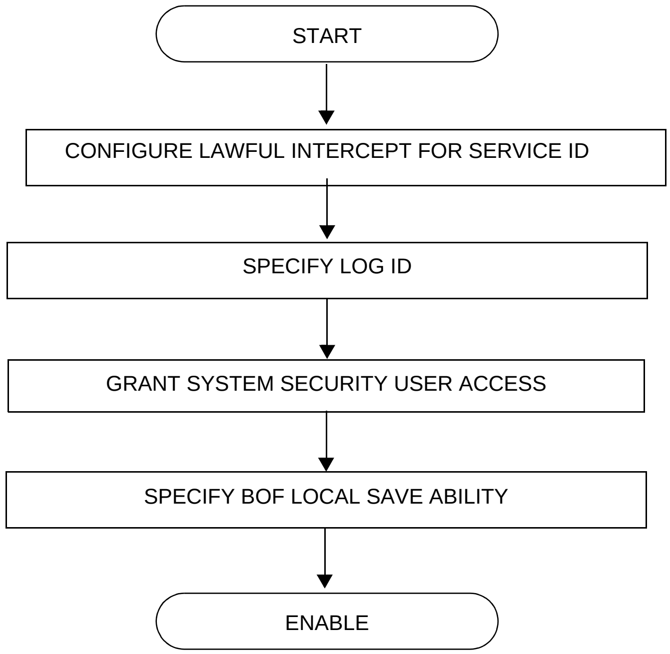

Figure 11: Lawful Intercept Configuration and Implementation Flow

Configuration Notes

This section describes mirroring configuration caveats.

- Multiple mirroring service IDs (mirror destinations) may be created within a single system.
- A mirrored source can only have one destination.
- 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.
- 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, 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.

Configuring Service Mirroring with CLI

This section provides information about service mirroring

Topics in this section include:

Mirror Configuration Overview

7750 SR-Series 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.

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 Alcatel-Lucent's implementation of mirroring, multiple source mirroring parameters can be specified to further identify traffic.

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

- Source IP address/mask
- Destination IP address/mask
- IP Protocol value
- Source port value/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

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.

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.

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 configured 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 displays an SNMP example showing 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 displays a user account configuration example.

```
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#
```

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 10 depicts a flow as to set an LI operator.

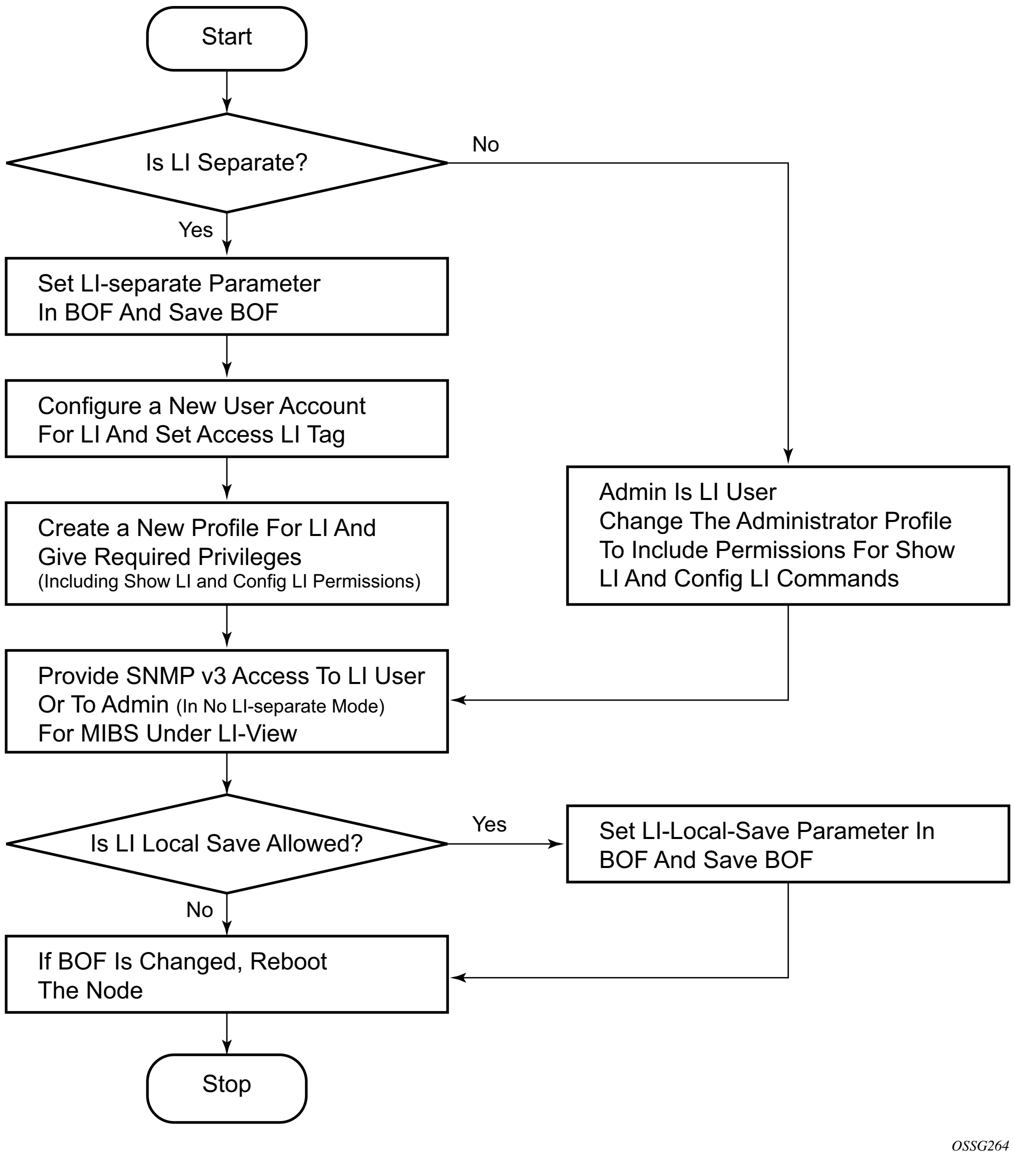

Figure 10: Creating an LI Operator Account

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 on page 49. 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 on page 49. 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)
- 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 on page 49.

- 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.

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 parametera 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.

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

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 displays a sample 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 displays the 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 displays a sample configuration of a remote mirrored service where the source is a port on ALA-A and the destination a SAP is on ALA-B.

```
*A:ALA-A>config>mirror# info
----------------------------------------------
        mirror-dest 1000 create
            sdp 2 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#
```

Mirror Classification Rules

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

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

Port

The `port` command associates a port to a mirror source. The port is identified by the port ID. The following displays the *port-id* syntax:

port-id:	slot/mda/port[.channel]		
	aps-*id*	**aps**-*group-id*[.*channel*]	
		aps	keyword
		group-id	1 — 64
	bundle-*type-slot/mda.bundle-num*		
		bundle	keyword
		type	ima
		bundle-num	1 — 128
	ccag-*id*	- **ccag**-*id.path-id*[*cc-type*]:*cc-id*	
		ccag	keyword
		id	1 — 8
		path-id	a, b
		cc-type	.sap-net, .net-sap
		cc-id	0 — 4094
	lag-id	1 — 64	
	egress	keyword	
	ingress	keyword	

The defined port can be 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 ATM, circuit-emulation (CEM), and PPP bundle groups cannot be used in a mirror source.

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

Table 2: Mirror Source Port Requirements

Port Type	Port Mode	Port Encap Type
faste/gige/xgige	access	dot1q, null
faste/gige/xgige	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`

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`

MAC filter

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

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

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

IP filter

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

Ingress mirrored packets are mirrored to the mirror destination prior to any ingress packet modifications. 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`* `…]`

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

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

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`

Subscriber

The subscriber command is used to add hosts of a subscriber to a mirroring service.

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

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 11) (within the same router) requires the following configurations:

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

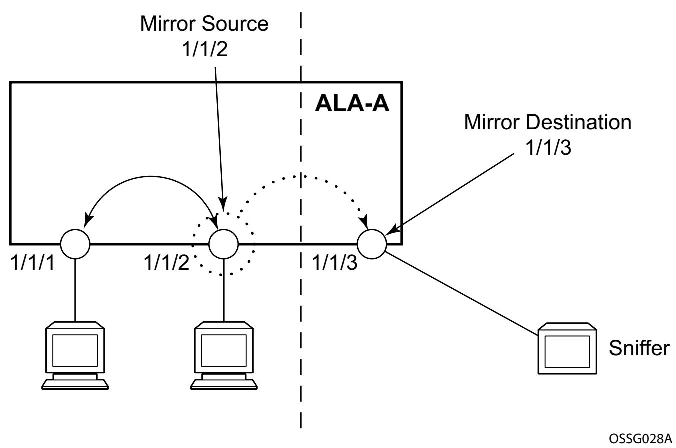

Figure 11: Local Mirrored Service Tasks

Each remote mirrored service (Figure 12) (across the network core) requires the following configurations:

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

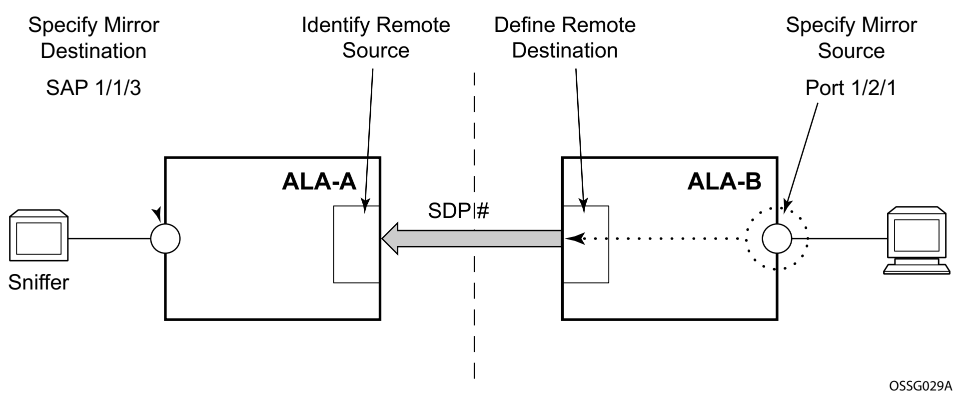

Figure 12: Remote Mirrored Service Configuration Example

Configuring a Local Mirror Service

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

The **mirror-source** commands are used as traffic selection criteria to identify traffic to be mirrored at the source. Each of these criteria are independent. For example, use the **debug>mirror-source>port** {*port-id* | **lag** *lag-id*} {[**egress**] [**ingress**]} command and **debug>mirror-source ip-filter** *ip-filter-id* **entry** *entry-id* [*entry-id…*] command to capture (mirror) traffic that matches a specific IP filter entry and traffic ingressing and egressing a specific port. A filter must be applied to the SAP or interface if only specific packets are to be mirrored. Note that slice-size is not supported by CEM encap-types or IP-mirroring.

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. Use the following CLI syntax to configure mirror destination parameters:

CLI Syntax:
```
config>mirror mirror-dest service-id [type {ether|frame-re-
lay|ppp|ip-only|atm-sdu|satop-e1|satop-t1|cesopsn|cesopsn-cas}] [create]
          description string
          fc fc-name
          sap sap-id [create]
          slice-size bytes
          no shutdown
```

CLI Syntax:
```
debug# mirror-source service-id
  ip-filter ip-filter-id entry entry-id [entry-id …]
  ingress-label label [label …]
  mac-filter mac-filter-id entry entry-id [entry-id …]
  port {port-id|lag lag-id} {[egress][ingress]}
  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 shutdown
```

CLI Syntax:
```
config>li
  li-source service-id
     ip-filter ip-filter-id entry entry-id [entry-id …]
     mac-filter mac-filter-id entry entry-id [entry-id …]
     sap sap-id {[ingress] [egress]}
     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]}] {[in-
```

```
                gress] [egress]}
        no shutdown
```

The following output displays 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 displays the 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
```

Configuring SDPs

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 Subscriber Services chapter.

Consider the following SDP characteristics:

- Configure either GRE or MPLS 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.
- The SDP IP address must be a 7750 SR system IP address.
- 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:

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

To configure the return path SDP, perform the same steps on the far-end 7750 SR router.

1. Select an originating node.
2. Create an SDP ID.
3. Select an encapsulation type.
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. When you configure a distributed Epipe SAP, you must identify an SDP ID. 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 displays 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#
```

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.

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

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

CLI Syntax:
```
debug> mirror-source service-id
  ip-filter ip-filter-id entry entry-id [entry-id …]
  ingress-label label [label …]
  mac-filter mac-filter-id entry entry-id [entry-id …]
  port {port-id|lag lag-id} {[egress][ingress]}
  sap sap-id {[egress][ingress]}
  sdp sap-id:[vc-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 shutdown
```

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

CLI Syntax:
```
config>mirror#
  mirror-dest service-id [type {ether|frame-relay|ppp|ip-on-
  ly|atm-sdu|satop-e1|satop-t1|cesopsn|cesopsn-cas}]
     description string
     fc fc-name
     remote-source
        far-end ip-addr ing-svc-label ing-svc-label
     sap sap-id
     sdp sdp-id[:vc-id][egr-svc-label [label|tldp]
     no shutdown
     slice-size bytes
```

CLI Syntax:
```
config>li
  li-source service-id
     ip-filter ip-filter-id entry entry-id [entry-id …]
     mac-filter mac-filter-id entry entry-id [entry-id …]
     port {port-id|lag lag-id} {[egress][ingress]}
     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]}] {[in-
                    gress] [egress]}
                no shutdown
```

The following displays the mirror destination, which is on ALA-A, configuration for mirror service 1216. This configuration specifies that the mirrored traffic coming from the mirror source (10.10.0.91) is to be directed to SAP 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 7750 SR and the core network.

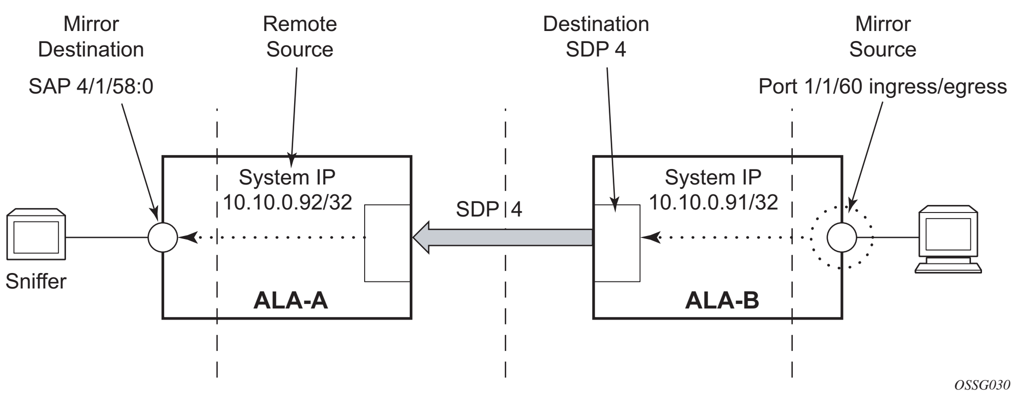

Figure 13: Remote Mirrored Service Tasks

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

The following displays the mirror destination configuration for mirror service 1216 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 displays 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
            sdp 4 egr-svc-label 5678
            slice-size 128
            no shutdown
        exit
---------------------------------------------
*A:ALA-B>config>mirror#
```

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

```
*A:ALA-A>config>service>sdp# info
---------------------------------------------
            description "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#
```

Configuring an ATM Mirror Service

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# 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
```

Configuring Lawful Intercept Parameters

The following display 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#
```

Pseudowire Redundancy for Mirror Services Configuration Example

A configuration based on Figure 14 is described.

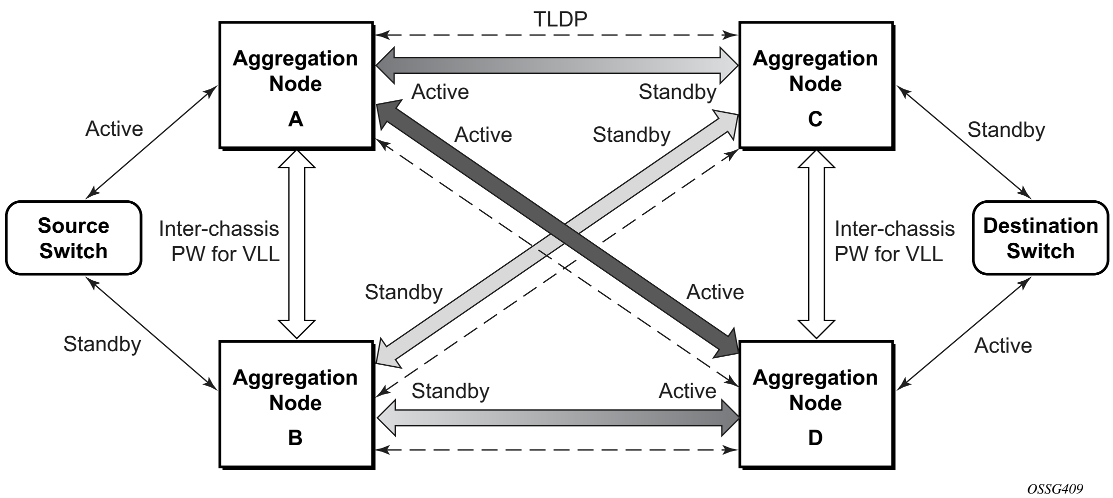

Figure 14: State Engine for Redundant Service to a Redundant Mirror Service

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.

```
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
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying a Local Mirrored Service on page 69
- Deleting a Local Mirrored Service on page 70
- Modifying a Remote Mirrored Service on page 71
- Deleting a Remote Mirrored Service on page 73

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

CLI Syntax:
```
config>mirror#
  mirror-dest service-id [type {ether|frame-relay|ppp|ip-on-
  ly|atm-sdu|atm-sdu|satop-e1|satop-t1|cesopsn|cesopsn-cas}]
     description description-string
     no description
     fc fc-name
     no fc
     remote-source
        far-end ip-address [ing-svc-label ing-svc-label|tldp]
        no far-end ip-address
     sap sap-id
     no sap
     sdp sdp-name [egr-svc-label egr-svc-label|tldp]
     no sdp
     [no] shutdown
```

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

CLI Syntax:
```
config>li
  li-source service-id
     ip-filter ip-filter-id entry entry-id [entry-id …]
     mac-filter mac-filter-id entry entry-id [entry-id …]
```

```
sap sap-id {[ingress] [egress]}
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]}] {[in-
      gress] [egress]}
no shutdown
```

Modifying a Local Mirrored Service

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

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

Example:

```
config>mirror# mirror-dest 103
       config>mirror>mirror-dest# shutdown
       config>mirror>mirror-dest# no sap
       config>mirror>mirror-dest# sap 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 displays 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#
```

Deleting a Local Mirrored Service

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

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

Example:ALA-A>`config>mirror# mirror-dest 103`

```
config>mirror>mirror-dest# shutdown
config>mirror>mirror-dest# exit
config>mirror# no mirror-dest 103
config>mirror# exit
```

Modifying a Remote Mirrored Service

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

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

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

Example:

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

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

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

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

```
*A:ALA-A>config>mirror# info
----------------------------------------------
        mirror-dest 104 create
            remote-source
                far-end 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
            sdp 4 egr-svc-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
```

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
```

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#
```

Mirror Service Command Reference

Command Hierarchies

Mirror Configuration Commands

config
- — **mirror**
  - — **mirror-dest** *service-id* [**type** *encap-type*] [create]
  - — **no mirror-dest** *service-id*
    - — **description** *description-string*
    - — **no description**
    - — [**no**] **enable-port-id**
    - — **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* [**ing-svc-label** *ing-vc-label* | **tldp**]
      - — **no far-end** *ip-address*
    - — **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**
        - — [**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*
    — **egress**
        — **vc-label** *egress-vc-label*
        — **no vc-label** [*egress-vc-label*]
    — **precedence** *precedence-value* | **primary**
    — **no precedence**
    — [**no**] **shutdown**
— [**no**] **shutdown**

Debug Commands

debug
— [**no**] **mirror-source** *mirror-dest-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* …]
    — **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**

Lawful Intercept Commands

config
— **li**
  — [**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* [*entry-id*...]
    — **no ip-filter** *ip-filter-id* [**entry** *entry-id*...]
    — **mac-filter** *mac-filter-id* **entry** *entry-id* [*entry-id*...]
    — **no mac-filter** *mac-filter-id* [**entry** *entry-id*...]
    — **nat**
      — [**no**] **classic-lsn-sub router** *router-instance* **ip** *ip-address*
        — **intercept-id** [1..4294967295]
        — **no intercept-id**
      — [**no**] **dslite-lsn-sub router router** *router-instance* **b4** *ipv6-prefix*
        — **intercept-id**[1..4294967295]
        — **no intercept-id**
      — **ethernet-header** [**da** *ieee-address*] [**sa** *ieee-address*] [**etype** *ethertype*]
      — **no ethernet-header**
      — [**no**] **ethernet-header** *sub-ident-string*
        — **intercept-id** [1..4294967295]
        — **no intercept-id**
      — [**no**] **l2-aware-sub** *sub-ident-string*
    — **sap** *sap-id* {[**ingress**] [**egress**]}
    — **no sap** *sap-id* [**ingress**] [**egress**]
    — [**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**]}
    — **no subscriber** *sub-ident-string*
  — **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*]
  — **save**

The following commands are also described in the 7750 SR OS Basic System Configuration Guide .

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

The following commands are also described in the 7750 SR OS System Management Configuration Guide.

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

Show Commands

show
— **debug** [*application*]
— **mirror mirror-dest** [*service-id*]
— **li**
— **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**]
— **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**

Configuration Commands

Generic Commands

description

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

Context config>mirror>mirror-dest
config>li>log>log-id

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.

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

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

shutdown

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

Context config>mirror>mirror-dest
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 source7750 SR-Series 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.

Mirror Destination Configuration Commands

far-end

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

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

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

The **far-end** command is used within the context of the **remote-source** node. It allows the definition of accepted remote sources for mirrored packets to this *mirror-dest-service-id*. Up to 50 **far-end** sources can be specified. If a far end router has not been specified, packets sent to the router are discarded.

The **far-end** command is used to define a remote source 7750 SR that may send mirrored packets to this 7750 SR for handling by this **mirror-dest** *service-id*.

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

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

Default No far end service ingress addresses are defined.

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

Values 1.0.0.1 — 223.255.255.254

The ingress service label must be manually defined using the **ing-svc-label** keyword. On the far end 7750 SR, the associated SDP **egr-svc-label** must be manually set and equal to the label defined in **ing-svc-label**.

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

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

The specified *ing-svc-label* must not have been used for any other service ID and must match the far end expected specific *egr-svc-label* for this 7750 SR. It must be within the range specified for manually configured service labels defined on this 7750 SR. It may be reused for other far end addresses on this *mirror-dest-service-id*.

Values 2048 — 18431

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

enable-port-id

Syntax [**no**] **enable-port-id**

Context configure>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.

Default no enable-port-id

endpoint

endpoint *endpoint-name* [**create**]
no endpoint *endpoint-name*

Context configure>mirror>mirror-dest
configure>mirror>mirror-dest>sap
configure>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 sdp's) 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
  - f 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 configure>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 *revert-time* — 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 — 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** *fc-name*
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 *fc-name*.

When the destination is on an SDP, the *fc-name* defines the DiffServ based egress queue that will be used to reach the destination. The *fc-name* also defines the encoded forwarding class of the encapsulation.

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

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

isa-aa-group

Syntax **isa-aa-group** *aa-group-id traffic-direction*

Context config>mirror>mirror-dest

Description This command specifies ISA AA group parameters.

Parameters *aa-group-id* — 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.

traffic-direction — 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** *service-id* [**type** *encap-type*] [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 7750 SR-Series router) or remotely, over the core of the network and have a far end 7750 SR-Series decode the mirror encapsulation.

The **mirror-dest** service is comprised of destination parameters that define where the mirrored packets are to be sent. It also specifies whether the defined *service-id* will receive mirrored packets from far end 7750 SR-Series 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 *service-id*. 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 *service-id* does not exist within the context of all defined services, the **mirror-dest** service is created and the context of the CLI is changed to that service ID. If the *service-id* exists within the context of defined **mirror-dest** services, the CLI context is changed for editing parameters on that service ID. If the *service-id* exists within the context of another service type, an error message is returned and CLI context is not changed from the current context.

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** *service-id* do not need to be removed or shutdown first. The **mirror-dest** *service-id* must be shutdown before the service ID can be removed. When the service ID is removed, all **mirror-source** 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 7750 SR-Series 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	*service-id*:	1 — 2147483647
	svc-name:	64 characters maximum

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

Values ether, frame-relay, ppp, ip-only, atm-sdu, satop-e1, satop-t1, cesopsn, cesopsn-cas

remote-source

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

Context config>mirror>mirror-dest

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

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

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

The **remote-source** node allows the source of mirrored packets to be on remote 7750 SR devices. The local 7750 SR will configure its network ports to forward packets associated with the *service-id* to the destination SAP. When **remote-source far-end** addresses are configured, an SDP is not allowed as a destination.

By default, the **remote-source** context contains no **far-end** addresses. When no **far-end** addresses have been specified, network remote devices will not be allowed to mirror packets to the local 7750 SR as a mirror destination. Packets received from unspecified **far-end** addresses will be discarded at network ingress.

The **no** form of the command restores the *service-id* to the default condition to not allow a remote 7750 SR access to the mirror destination. The **far-end** addresses are removed without warning.

Default No remote source devices defined

sap

Syntax
sap *sap-id* [**create**] [**no-endpoint**]
sap *sap-id* [**create**] **endpoint** *name*
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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 355 for command syntax.

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.

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** *milliseconds* [**payload-size** *bytes*]
packet payload-size *bytes*
no packet *bytes*

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..4	16
	N = 5..15	8
	N >= 16	5
nxDS0WithCas (E1)	N	8
nxDS0WithCas (T1)	N	12

Parameters *milliseconds* — 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..4	N x 32
	N = 5..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 — 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. Zeroes 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 *policy-id* — The QoS policy ID to associate with SAP for the mirrored service. The policy ID must already exist.

Values 1 — 65535

service-name

Syntax **service-name** *service-name*
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** *bytes*
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 *bytes* — The number of bytes to which mirrored frames will be truncated, expressed as a decimal integer.

Values 128 — 9216

spoke-sdp

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

Context config>mirror>mirror-dest

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

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

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

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

One method is the IP Generic Router Encapsulation (GRE) encapsulation, which has no state in the core of the network. GRE does not specify a specific path to a router. A GRE-based SDP uses the underlying IGP routing table to find the best next hop to the far end router.

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

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

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

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

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

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

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

Values 1 — 17407

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.

egress

Syntax **egress**

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

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

vc-label

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

Context config>mirror>mirror-dest>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 — 1048575

precedence

precedence *precedence-value* **| 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.

Context An SDP is created with type secondary and with the lowest precedence value of 4.

Parameters *prec-value* — The precedence of the SDP.

Values 1-4

primary — A special value of the precedence which assigns the SDP the lowest precedence and enables the revertive behavior.

Mirror Source Configuration 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 mirrored packet is properly referenced by multiple mirror sources (for example, a SAP on one **mirror-source** and a port on another **mirror-source**), then the packet is mirrored to a single *mirror-dest-service-id* based on the following hierarchy:

1. Filter entry
2. Subscriber mirror priority
3. Service access port (SAP)
4. Physical port

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

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

By default, all **mirror-dest** service IDs have a **mirror-source** associated with them. The **mirror-source** is not technically created with this command. Instead the service ID provides a contextual node for storing the current mirroring sources for the associated **mirror-dest** service ID. The **mirror-source** is created for the mirror service when the operator enters the **debug>mirror-source** *svcId* for the first time. 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 *service-id* — The mirror destination service ID for which match criteria will be defined. The *service-id* must already exist within the system.

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

ip-filter

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

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 *mirror-dest-service-id* of the **mirror-source**.

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

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

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

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

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

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

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

Default IP filter mirroring is not defined.

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

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

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

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

mac-filter

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

Context debug>mirror-source

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

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

The MAC filter must already exist in order for the command to execute. Filters are configured in the config>filter context. If the MAC filter does not exist, an error will occur. If the filter exists but has not been associated with a SAP or IP interface, an error is not 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 *mac-filter-id* — The MAC filter ID whose entries are mirrored. If the *mac-filter-id* does not exist, an error will occur and the command will not execute. Mirroring of packets will commence once the *mac-filter-id* is defined on a SAP.

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

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

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

port

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

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.

Syntax: port-id: *slot/mda/port*[.*channel*]
bundle-id:bundle-*type-slot*/*mda*.*bundle-num*
bundle keyword
type ima, fr, ppp
bundle-num 1 — 336
bpgrp-id:bpgrp-*type-bpgrp-num*
bpgrp keyword
type ima, ppp
bpgrp-num 1 — 2000
aps-id: aps-*group-id*.*channel*
aps keyword
group-id 1 — 64

ccag-id:	ccag-*id*.*path-id* *cc-type*:cc-id	
	ccag	keyword
	id	1 — 8
	path-id	a, b
	cc-type	.sap-net, .net-sap
	cc-id	0 — 4094

ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
	ccag keyword	
	id 1 — 8	
	path-id	a, b
	cc-type	.sap-net, .net-sap
	cc-id	0 — 4094

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

Values 1 — 200

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** *sap-id* {[**egress**] [**ingress**]}
no sap *sap-id* [**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 *sap-id* is owned by the service ID of the service in which it was created. The SAP is only referenced in the mirror source name for mirroring purposes. The mirror source association does not need to be removed before deleting the SAP from its service ID. If the SAP is deleted from its service ID, the mirror association is removed from the mirror source.

More than one SAP can be associated within a single **mirror-source**. Each SAP has its own **ingress** 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.

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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 355 for command syntax.

channel-id — The SONET/SDH or TDM channel on the port of the SAP. A period separates the physical port from the *channel-id*. The port must be configured as an access port.

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.

ingress-label

Syntax [**no**] **ingress-label** *label* [*label …*up to 8 max]
no ingress-label *label* [*label …*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 7750 SR 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 *label* — 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 — 1048575. The local MPLS stack may not support portions of this range.

Lawful Intercept Commands

li

Syntax **li**

Context config

Description This command configures the context to configure lawful intercept (LI) parameters.

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.

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

Default **li-filter-lock-state locked**

Parameters **locked —** When an Lawful Interface 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 all users even when an li-source references an entry in that filter.

unlocked-for-all-users — Filters can continue to be edited by LI users only even when an li-source references an entry in that filter.

li-source

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

Context config>li

Description This command configures a lawful intercept (LI) mirror source.

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.

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

ip-filter

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

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* [*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.

mac-filter

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

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 7750 SR OS 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.

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 *mac-filter-id* — Specifies the MAC filter ID. If the *mac-filter-id* does not exist, an error will occur and the command will not execute.

entry *entry-id* [*entry-id* …] **—** The MAC filter entries to use as match criteria.

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

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

Default Base

ip *ip-address* **—** Specifies the IP address in a.b.c.d format.

intercept-id

Syntax **intercept-id** [1..4294967295]
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

Description This command configures the intercept identifier.

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

Parameters **1..4294967295 —** The intercept identifier range.

Values 1..4294967295

dslite-lsn-sub router

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:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x - [0..FFFF]H d - [0..255]D
<length>	: [0..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 *sub-ident-string* — Specifies a source name.

sap

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

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.

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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 355 for command syntax.

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** *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*

Context config>li>li-source

Description This command adds hosts of a subscriber to mirroring service.

Parameters *sub-ident-string* — Specifies the name of the subscriber identification policy.

sap-id — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 355 for command syntax.

ip-address — The service IP address (system IP address) of the remote 7750 SR device sending LI traffic. If 0.0.0.0 is specified, any remote 7750 SR is allowed to send to this service.

Values 1.0.0.1 — 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 — Specifies the SLA profile name.

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.

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 *log-id* is used to direct events, alarms/traps, and debug information to respective destinations.

Parameters *log-id* — The log ID number, expressed as a decimal integer.

Values 1 — 100

filter

Syntax **filter** *filter-id*
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 *log-id*.

Default **no filter — No event filter policy is specified for a** *log-id.*

Parameters *filter-id* — The event filter policy ID is used to associate the filter with the *log-id* configuration. The event filter policy ID must already be defined in **config>log>filter** *filter-id*.

Values 1 — 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 requestor 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 *size* — The size parameter indicates the number of events that can be stored into memory.

Default 100

Values 50 — 1024

save

Syntax **save**

Context config>li

Description This command is required to save LI configuration parameters.

Other LI Configuration Commands

The following commands are also described in the 7750 SR OS Basic System Configuration 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.

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 *user-profile-name* can consist of up to 32 alphanumeric characters.

The **no** form of the command deletes a user profile.

Default **user-profile default**

Parameters *user-profile-name* — 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

Show Commands

debug

Syntax **debug** [*application*]

Context **show**

Description This command displays set debug points.

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

Values: service, ip, ospf, ospf3, bgp, mtrace, rip, isis, mpls, rsvp, ldp, mirror, vrrp, system, filter, subscriber-mgmt, radius, lag, oam, frame-relay, local-dhcp-server, ,mld, pim

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 l2 l1 h1 nc ingress
        subscriber "user11" sap 1/1/2:1 mac 00:00:01:00:00:02 fc be l1 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 l1 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:alu1# 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:alu1#
```

active-subscribers

Syntax **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*]

Context show>service

Description This command displays active subscriber information.

Parameters *sub-ident-string* — Specifies an existing subscriber identification string.

sap *sap-id* — Specifies the physical port identifier portion of the SAP definition. See “Common CLI Command Descriptions” on page 355 for command syntax.

sla-profile-name — Displays an existing SLA profile name.

hierarchy — Displays the subscriber hierarchy.

summary — Displays subscriber summary.

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   l2 af l1 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   l2 af l1 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   l2 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 l2 af l1 h2 ef h1 nc
                    Egress  mirror:      100   be l2 af l1 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   l1 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   l2 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   l1 nc
                    Egress  mirror:      100   l1 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 l2 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 l2 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 l1 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 l1 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 l2 af l1 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 l2 af l1 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 l2 af l1 h2 ef h1 nc
                    Egress  mirror:      100   be l2 af l1 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 l2 af l1 h2 ef h1 nc
                    Egress  mirror:      100   be l2 af l1 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 l2 af l1 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 **Show Service-Using Mirror —** The following table describes service-using mirror output fields:

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

Sample Output

```
A:ALA-48# show service service-using mirror
===============================================================================
Services [mirror]
===============================================================================
ServiceId     Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
218           Mirror    Up     Down       1                 04/08/2007 13:49:57
318           Mirror    Down   Down       1                 04/08/2007 13:49:57
319           Mirror    Up     Down       1                 04/08/2007 13:49:57
320           Mirror    Up     Down       1                 04/08/2007 13:49:57
1000          Mirror    Down   Down       1                 04/08/2007 13:49:57
1216          Mirror    Up     Down       1                 04/08/2007 13:49:57
1412412       Mirror    Down   Down       1                 04/08/2007 13:49:57
-------------------------------------------------------------------------------
Matching Services : 7
===============================================================================
A:ALA-48#
```

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 — 2147483647

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#
```

log

Syntax **log**

Context show>li

Description Displays Lawful Intercept event log information.

log-id

Syntax **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**]

Context show>li>log

Description Displays information for specified log.

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

Values 1 — 100

severity-level — Specifies the severity level.

Values cleared, indeterminate, critical, major, minor, warning

application — Specifies the application name.

Values application_assurance, aps, atm, bgp, cflowd, chassis, debug, dhcp, dhcps, dot1ag, efm_oam, filter, gsmp, igmp, igmp_snooping, ip, isis, lag, ldp, li, logger, mc_redundancy, mirror, mld, mld_snooping, mpls, msdp, ntp, oam, ospf, pim, pim_snooping, port, ppp, pppoe, rip, route_policy, rsvp, security, snmp, stp, svcmgr, system, user, vrrp, vrtr

from-seq [*to-seq*] — Specifies the sequence value.

Values 1 — 4294967295

count — Specifies the count.

Values 1 — 4294967295

subject — Specifies a subject string to match.

reexp — Specifies to use a regular expression match.

ascending|descending — Specifies the sort direction

router-instance — Specifies the router instance.

status

Syntax **status**

Context show>li

Description Displays Lawful Intercept status information.

Sample Output

```
*A:sim138# 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:sim138#
```

li

Syntax **li**

Context show

Description Displays Lawful Intercept (LI) information.

mirror mirror-dest

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

Context show

Description This command displays mirror configuration and operation information.

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

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

Label	Description
`Service Id`	The service ID associated with this mirror destination.
`Type`	Entries in this table have an implied storage type of “volatile”. The configured mirror source information is not persistent.
`Admin State`	`Up` — The mirror destination is administratively enabled.
	`Down` — The mirror destination is administratively disabled.
`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.

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             : 0                 Enable 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#
```

Debug Commands

subscriber

Syntax **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*

Context debug>mirroring-source

Description This command adds hosts of a subscriber to mirroring service.

Parameters *sub-ident-string* — Specifies the name of the subscriber identification policy.

sap-id — Specifies the physical port identifier portion of the SAP definition. See “Common CLI Command Descriptions” on page 355 for command syntax.

ip *ip-address* — The service IP address (system IP address) of the remote 7750 SR device sending LI traffic.

Values 1.0.0.1 — 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 *sla-profile-name* — 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 **ingress-label l***abel* [*label …*up to 8 max]
no ingress-label [*label* [*label …*up to 8 max]]

Context debug>mirror-source

Description This command configures mirroring of ingress MPLS frames with a specific MPLS label to a mirror destination.

OAM and SAA

In This Chapter

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

Topics in this chapter include:

OAM Overview

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.

Two-Way Active Measurement Protocol

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

There are four logical entities in TWAMP: the control-client, the session-sender, the server, 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") with which the two-way measurements are being performed. 7750 SRacts as the server.

The control-client and server establish a TCP connection and exchange TWAMP-Control messages over this connection. When the control-client wants to start testing, the client communicates the test parameters to the server. If the server agrees to conduct the described tests, the test begin as soon as the client sends a Start-Sessions message. As part of a test, the session-sender sends a stream of UDP-based test packets to the session-reflector, and the session reflector responds to each received packet with a response UDP-based test packet. When the session-sender receives the response packets from the session-reflector, the information is used to calculate two-way delay, packet loss, and packet delay variation between the two devices.

LSP Diagnostics

The 7750 SR LSP diagnostics are implementations of LSP ping and LSP traceroute based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. In an LDP ECMP network, a unique-path trace can be accomplished by specifying a unique 127/8 IP address for the **path-destination** *ip-address* parameter. Note that the 7750 SR can send multipath type 0 or 8, and up to a maximum of 36 bytes for multipath length (refer to RFC 4379 for more details). The 7750 SR supports unique-path trace on an LER of an LDP ECMP path. LSP ping, as described in the draft, provides a mechanism to detect dataplane failures in MPLS LSPs. LSP ping and LSP traceroute are modeled after the ICMP echo request/reply used by ping and traceroute to detect and localize faults in IP networks.

For a given FEC, LSP ping verifies whether the packet reaches the egress label edge router (LER), while in LSP traceroute 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.

LSP Ping for RSVP P2MP LSP (P2MP)

Note: For more information about P2MP refer to the 7750 SR OS MPLS Guide.

The P2MP LSP ping complies to draft-ietf-mpls-p2mp-lsp-ping-06, *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 7750 nodes, they will be able to parse the list of egress LER addresses and will reply. Note however that draft-ietf-mpls-p2mp-lsp-ping-06 specifies that only the top address in the P2MP egress identifier TLV must be inspected by an egress LER. When interoperating with other implementations, an 7750 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.

A 7750 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, a 7750 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 7750 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 however 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 a 7750 ingress LER node receives the new multipath type field with the list of egress LER addresses in an echo reply message from another vendor implementation, it will ignore but will not cause an error in processing the downstream mapping TLV.

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

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

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

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

LSP Trace for RSVP P2MP LSP

The P2MP LSP trace complies to draft-ietf-mpls-p2mp-lsp-ping-06.

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 downstream mapping TLV is modified when used over a P2MP LSP (draft-ietf-mpls-p2mp-lsp-ping-06):

- A new B-flag is added to the downstream mapping TLV to indicate that the reporting LSR is not a branch for this LSP (cleared to zero) or is a branch (set to one).
- A new E-flag is added to the downstream mapping TLV to indicate that the reporting LSR is not a bud node for this LSP (cleared to zero) or is a bud node (set to one).
- The flags are placed in the fourth byte of the TLV that was previously reserved as shown below. All other fields are unchanged from their definitions in RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*, except for the additional information that can be carried in the multipath information.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               MTU             | Address Type  | Reserved  |E|B|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Downstream IP Address (4 or 16 octets)            |
```

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Downstream Interface Address (4 or 16 octets)         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Hash Key Type | Depth Limit   |        Multipath Length       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
.                                                               .
.                     (Multipath Information)                   .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Downstream Label                |    Protocol   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
.                                                               .
.                                                               .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Downstream Label                |    Protocol   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 15: Modifications to the Downstream Mapping TLV

Note that the E-flag and B-flag are at bit positions 28 and 29 respectively. Bits 30 and 31 are already used as per RFC 4379. This error was reported to the authors of the draft.

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 responds, it sets the B-flag in the downstream mapping TLV to indicate to the sender of the echo request message it has other branches for this LSP. A bud LSR will also set the E-flag in the downstream mapping TLV to indicate to the sender of the echo request message that it is also an egress LER for the P2MP LSP when the traced egress is reachable via a downstream branch. In this case, the return code must correspond to the LSR role and must code #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.

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

When a 7750 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)

   A
 /  \
B -- C
|
D
| \
F  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 **DSMAP mismatch** error because the DSMAP 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 "DSMAP mismatch" 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 DSMAP. This time node D will respond with a success and will include its downstream DSMAP 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.

A 7750 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 7750 ingress LER behavior is to always proceed to the next ttl probe when it receives an OK response to a probe or when it times out on a probe. If however it receives replies with an error return code, it must wait until it receives an OK response or it times out. If it times out without receiving an OK reply, the LSP trace must be aborted.

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

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

SDP Diagnostics

The 7750 SR OS SDP diagnostics are SDP ping and SDP MTU path discovery.

SDP Ping

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

For a uni-directional test, SDP ping tests:

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

For a round-trip test, SDP ping uses a local egress SDP ID and an expected remote SDP ID. Since SDPs are uni-directional tunnels, the remote SDP ID must be specified and must exist as a configured SDP ID on the far-end 7750 SR.7750 SR OS MG SDP round trip testing is an extension of SDP connectivity testing with the additional ability to test:

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

SDP MTU Path Discovery

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

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

Service Diagnostics

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

Service ping is initiated from a 7750 SR router to verify round-trip connectivity and delay to the far-end of the service. Alcatel-Lucent'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

VPLS MAC Diagnostics

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

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

The VPLS OAM tools are:

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

MAC Ping

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

In the control plane, a MAC ping is forwarded along the flooding domain if no MAC address bindings exist. If MAC address bindings exist, then the packet is forwarded along those paths (if they are active). Finally, a response is generated only when there is an egress SAP binding to that MAC address. A control plane request is responded to via a control reply only.

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

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

MAC Trace

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

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

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

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

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

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

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

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

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

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

CPE Ping

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

The **cpe-ping** command extends this capability to detecting end-station IP addresses inside a VPLS. A CPE ping for a specific destination IP address within a VPLS will be translated to a MAC-ping towards a broadcast MAC address. Upon receiving such a MAC ping, each peer PE within the VPLS context will trigger an ARP request for the specific IP address. The PE receiving a response to this ARP request will report back to the requesting 7750 SR. It is encouraged to use the source IP address of 0.0.0.0 to prevent the provider's IP address of being learned by the CE.

MAC Populate

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

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

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

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

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

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

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

MAC Purge

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

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

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

VLL Diagnostics

VCCV Ping

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

VCCV-Ping Application

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

1. Use of a Router Alert Label immediately above the VC label. This method has the drawback that if ECMP is applied to the outer LSP label (for example, transport label), the VCCV message will not follow the same path as the user packets. This effectively means it will not troubleshoot the appropriate path. This method is supported by the 7750 SR.
2. Use of the OAM control word as illustrated in Figure 15.

```
 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          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 16: OAM Control Word Format

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

The use of the OAM control word assumes that the draft-martini control word is also used on the user packets. This means that if the control word is optional for a VLL and is not configured, the 7750 SR PE node will only advertise the router alert label as the CC capability in the Label Mapping message. This method is supported by the 7750 SR.

3. Set the TTL in the VC label to 1 to force PE2 control plane to process the VCCV message. This method is not guaranteed to work under all circumstances. For instance, the draft mentions some implementations of penultimate hop popping overwrite the TTL field. This method is not supported by the 7750 SR.

When sending the label mapping message for the VLL, PE1 and PE2 must indicate which of the above OAM packet encapsulation methods (for example, which control channel type) they support. This is accomplished by including an optional VCCV TLV in the pseudowire FEC Interface Parameter field. The format of the VCCV TLV is shown in Figure 16.

```
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    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 17: VCCV TLV

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

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

- 0x00 None of the following VCCV control channel types are supported
- 0x01 PWE3 OAM control word (see Figure 15)
- 0x02 MPLS Router Alert Label
- 0x04 MPLS inner label TTL = 1

If both PE nodes support more than one of the CC types, then a 7750 SR 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 7750 SR.

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 7750 SR.

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 7750 SR.
2. Reply via an IPv4/IPv6 UDP packet. This mode is supported by the 7750 SR.
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 7750 SR.
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 7750 SR.

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 7750 SR 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 7750 SR nodes. The VCCV ping feature can test connectivity of a VLL with any third party node which is compliant to RFC 5085.

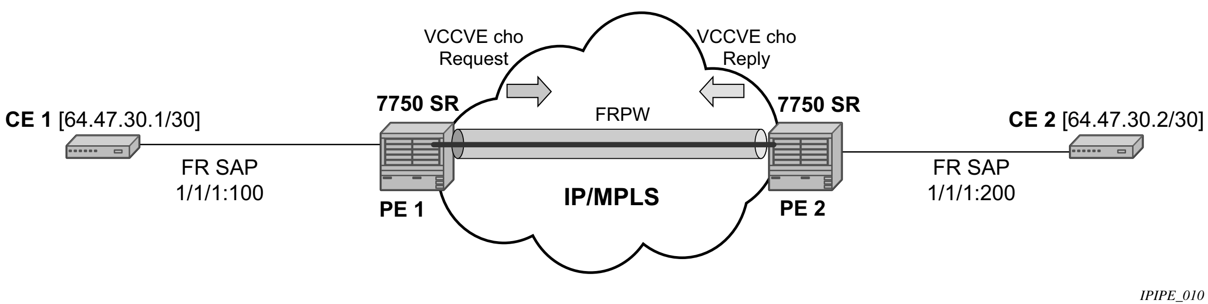

Figure 18: VCCV-Ping Application

VCCV-Ping in a Multi-Segment Pseudowire

Figure 21 displays and 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 functions:

1. VCCV ping to a destination PE. A VLL FEC Ping is a message sent by T-PE1 to test the FEC at T-PE2. The operation at T-PE1 and T-PE2 is the same as in the case of a single-segment pseudowire. The pseudowire switching node, S-PE1, pops the outer label, swaps the inner (VC) label, decrements the TTL of the VC label, and pushes a new outer label. The 7750 SR 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.

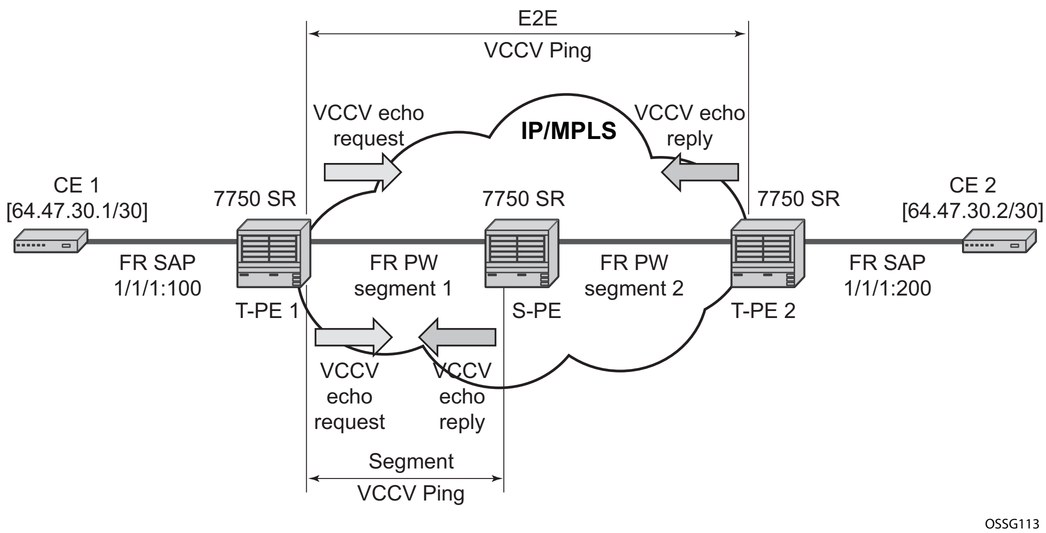

Figure 19: VCCV-Ping over a Multi-Segment Pseudowire

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 on page 134. 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.

VCCV for Static Pseudowire Segments

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

Detailed VCCV-Trace Operation

In Figure 21 on page 148 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.

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

Sending a VCCV Echo Request

When in the ping mode of operation, the sender of the echo request message requires the FEC of the last segment to the target S-PE/T-PE node. This information can either be configured manually or be obtained by inspecting the corresponding sub-TLV's of the pseudowire switching point TLV. However, the pseudowire switching point TLV is optional and there is no guarantee that all S-PE nodes will populate it with their system address and the pseudowire-id of the last pseudowire segment traversed by the label mapping message. Thus the 7750 SR 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.

Receiving an VCCV Echo Request

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

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

Receiving an VCCV Echo Reply

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

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

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.

IGMP Snooping Diagnostics

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.

ATM Diagnostics

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 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).

End-to-End Testing of Paths in an LDP ECMP Network

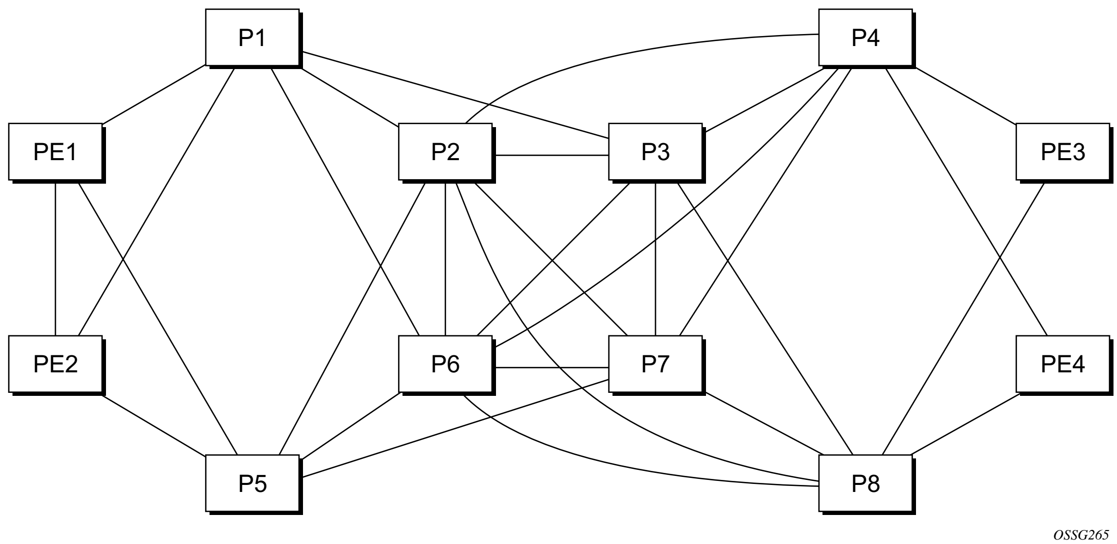

Figure 20: Network Resilience Using LDP ECMP

Figure 19 depicts 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.

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 ECMP OAM 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 NLHFE entry can also result from a corruption in the control plane at that node.

LDP ECMP Tree Building

The 7750 SR 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 7750 SR 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 7750 SR 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 7750 SR 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*:

```
PE1 ---- A ----- B ----- C ------ G ----- H ---- PE2
          \       \---- D ------/       /
               \      \--- E------/       /
                -- F --------------------/
```

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 7750 SR 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 7750 SR 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.

Periodic Path Exercising

The periodic path exercising 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 7750 SR 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.

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
1DM	One way Delay Measurement (Y.1731)
AIS	Alarm Indication Signal
CCM	Continuity check message
CFM	Connectivity fault management
DMM	Delay Measurement Message (Y.1731)
DMR	Delay Measurement Reply (Y.1731)
LBM	Loopback message
LBR	Loopback reply
LTM	Linktrace message
LTR	Linktrace reply

Acronym	Callout (Continued)
ME	Maintenance entity
MA	Maintenance association
MA-ID	Maintenance association identifier
MD	Maintenance domain
MEP	Maintenance association end point
MEP-ID	Maintenance association end point identifier
MHF	MIP half function
MIP	Maintenance domain intermediate point
OpCode	Operational Code
RDI	Remote Defect Indication
TST	Ethernet Test (Y.1731)
SLM	Synthetic Loss Message
SLR	Synthetic Loss Reply (Y.1731)

ETH-CFM Building Blocks

The IEEE and the ITU-T use their own nomenclature when describing administrative contexts and functions. This introduces a level of complexity to configuration, discussion and different vendors naming conventions. The SR and ESS 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) form the ASCII table

3 (2-octet integer) — 0 — 65535

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

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

Note: When a VID is used as the short MA name, 802.1ag will not support VLAN translation because the MA-ID must match all the MEPs. The default format for a short MA name is an integer. Integer value 0 means the MA is not attached to a VID. This is useful for VPLS services on SR/ESS 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
configuer eth-cfm domain 4 format string name IEEE-Domain level 4

show eth-cfm domain
===============================================================================
CFM Domain Table
===============================================================================
Md-index   Level Name                                       Format
-------------------------------------------------------------------------------
3          3                                                none
4          4     IEEE-Domain                                charString
===============================================================================
```

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
----------------------------------------------
*A:cses-E01>config>eth-cfm#  show eth-cfm association

===============================================================================
CFM Association Table
===============================================================================
Md-index   Ma-index   Name                     CCM-intrvl Hold-time Bridge-id
-------------------------------------------------------------------------------
3          1          123456789abcd            10         n/a       100
3          2          Y1731ContextIEEEFormat   10         n/a       300
4          1          UpTo45CharactersForIEEE* 1          n/a       100
===============================================================================
```

* indicates that the corresponding row element may have been truncated..

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 this chapter.

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

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

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

This is a general table that indicates the 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 3: 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					No
	SAP	Yes	Yes	Yes	-
	Spoke-SDP	Yes	Yes	Yes	-
	Mesh-SDP	Yes	Yes	No	-
B-VPLS					Yes
	SAP	Yes	Yes	Yes	-
	Spoke-SDP	Yes	Yes	Yes	-
	Mesh-SDP	Yes	Yes	No	-
I-VPLS					No
	SAP	Yes	Yes	Yes	-
	Spoke-SDP	Yes	Yes	Yes	-
	Mesh-SDP	Yes	Yes	No	-
M-VPLS					No

Service	Ethernet Connection	Down MEP	Up MEP	MIP	Virtual MEP
	SAP	Yes	Yes	Yes	-
	Spoke-SDP	Yes	Yes	Yes	-
	Mesh-SDP	Yes	Yes	No	-
PBB EPIPE					No
	SAP	Yes	Yes	Yes	-
	Spoke-SDP	Yes	Yes	Yes	-
	Mesh-SDP	Yes	Yes	No	-
IPIPE					No
	SAP	Yes	No	No	-
	Ethernet-Tunnel SAP	Yes	No	No	-
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	-
Note1	Ethernet-Tunnel (Control) SAP	Yes	No	No	-
	Ethernet-Tunnel (Path/Member)	Yes	Yes	No	-
	Ethernet-Ring (Data)	Yes	No	No	-

Note1: Ethernet-Tunnels and Ethernet-Rings are not configurable under all service types. Any service restrictions for MEP direction or MIP support will override the generic capability of the thernet-Tunnel or Ethernet-Ring MPs. Please check the applicable user guide for applicability.

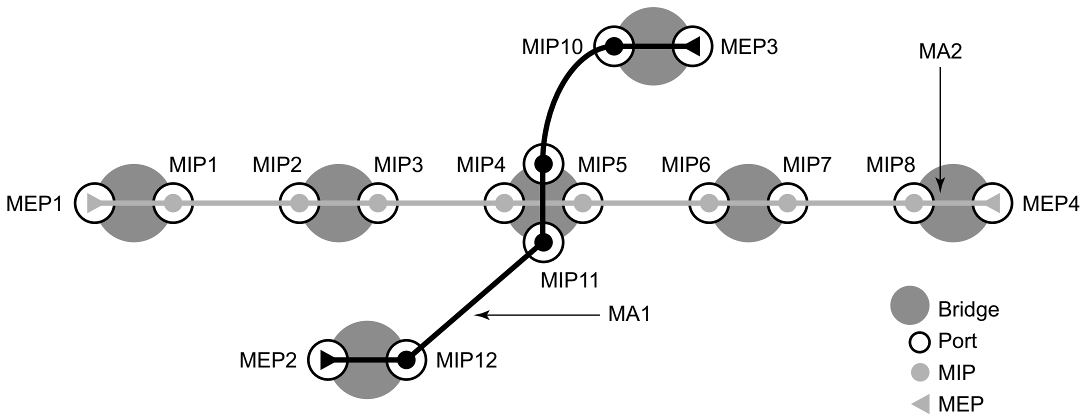

Figure 21: MEP and MIP

Figure 21 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.

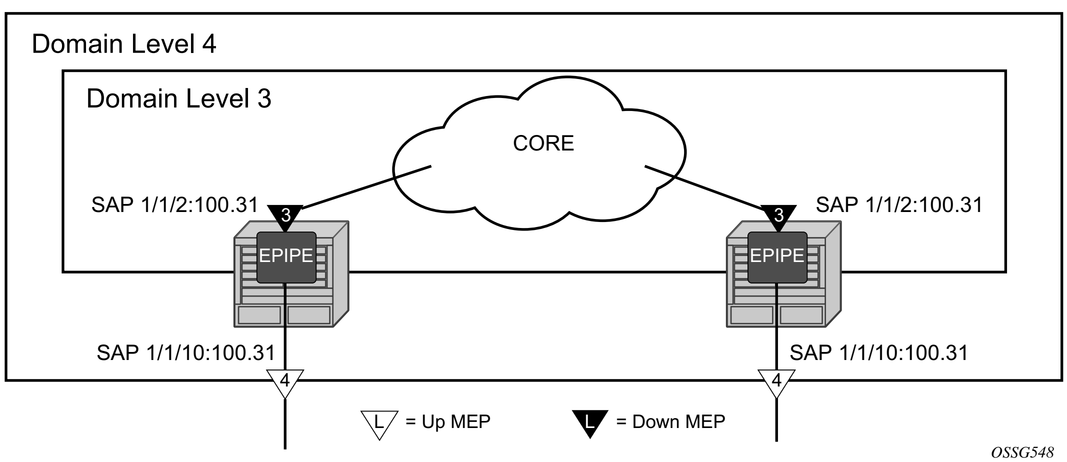

Figure 22: MEP Creation

```
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

*A:cses-E01>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
----------------------------------------------
*A:cses-E02>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
----------------------------------------------
*A:cses-E02>config>service>epipe#
```

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 cas,e the Level 3 domain is completely contained in a Level 4 domain.

The following display was taken from NODE1.

```
show eth-cfm cfm-stack-table
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM SAP Stack Table
===============================================================================
Sap                Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/1/2:100.31         3 Down          3          1  111 90:f3:01:01:00:02 ------
1/1/10:100.31        4   Up          4          1  101 d0:0d:1e:00:01:01 ------
===============================================================================
```

Figure 22 illustrates the creation of and explicit MIP.

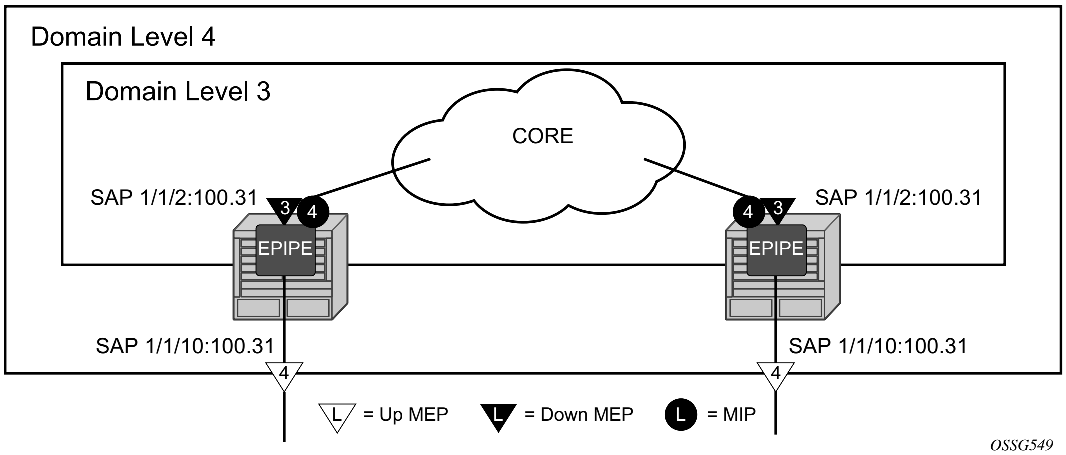

Figure 23: MIP Creation Example (NODE1)

```
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
            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 has been included. 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.

The following output is from Node 1.

```
show eth-cfm cfm-stack-table
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM SAP Stack Table
===============================================================================
Sap                 Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/1/2:100.31          3 Down          3          1  111 d0:0d:1e:00:01:11 ------
1/1/2:100.31          4 Both          4          2  MIP 90:f3:01:01:00:02 ------
1/1/10:100.31         4   Up          4          1  101 d0:0d:1e:00:01:01 ------
===============================================================================
```

Figure 23 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.

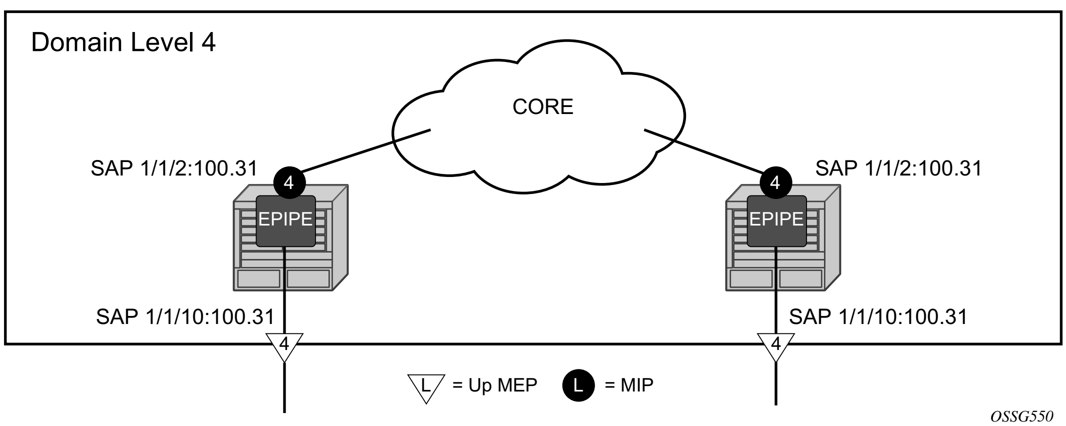

Figure 24: MIP Creation Default

```
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
----------------------------------------------

# show eth-cfm cfm-stack-table
================================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx
================================================================================
CFM SAP Stack Table
================================================================================
Sap                 Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
--------------------------------------------------------------------------------
1/1/2:100.31          4 Both          4          2  MIP d0:0d:1e:01:01:01 ------
1/1/10:100.31         4   Up          4          1  101 d0:0d:1e:00:01:01 ------
================================================================================

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
----------------------------------------------
```

```
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
            no shutdown
----------------------------------------------

# show eth-cfm cfm-stack-table
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM SAP Stack Table
===============================================================================
Sap                Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/1/2:100.31         4 Both          4          2  MIP d0:0d:1e:01:01:02 ------
1/1/10:100.31        4   Up          4          1  102 d0:0d:1e:00:01:02 ------
===============================================================================
```

Figure 24 shows the detailed IEEE representation of MEPs, MIPs, levels and associations, using the standards defined icons.

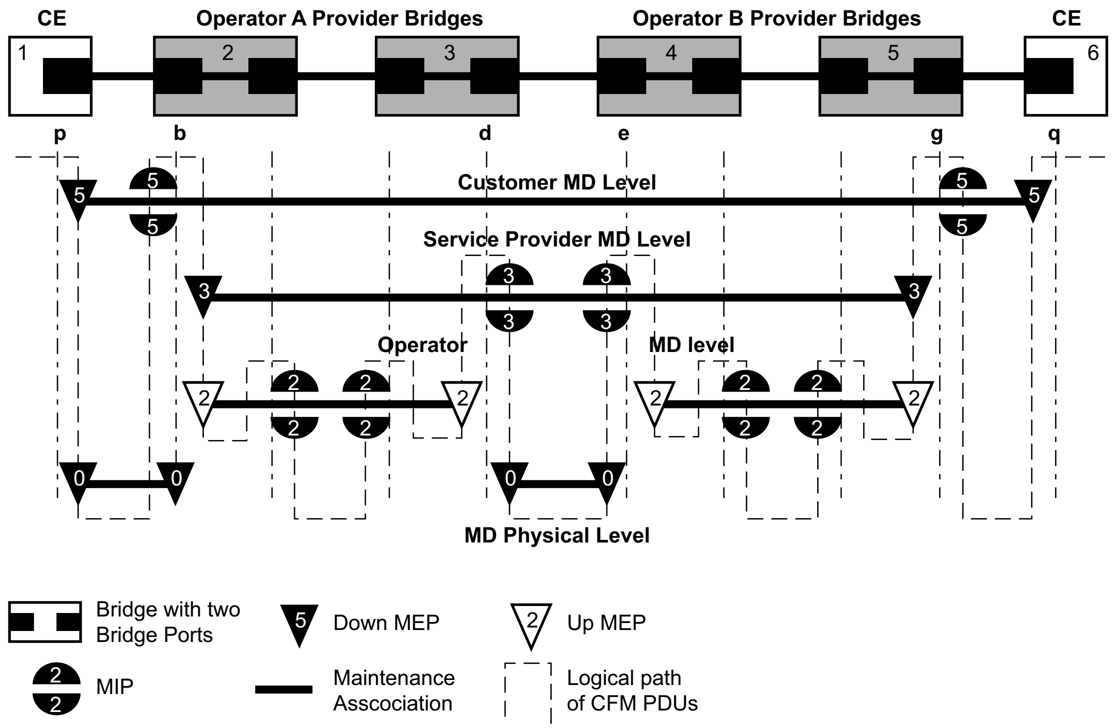

Figure 25: MEP, MIP and MD Levels

Loopback

A loopback message is generated by an MEP to its peer MEP or a MIP (Figure 25). The functions are similar to an IP ping to verify Ethernet connectivity between the nodes.

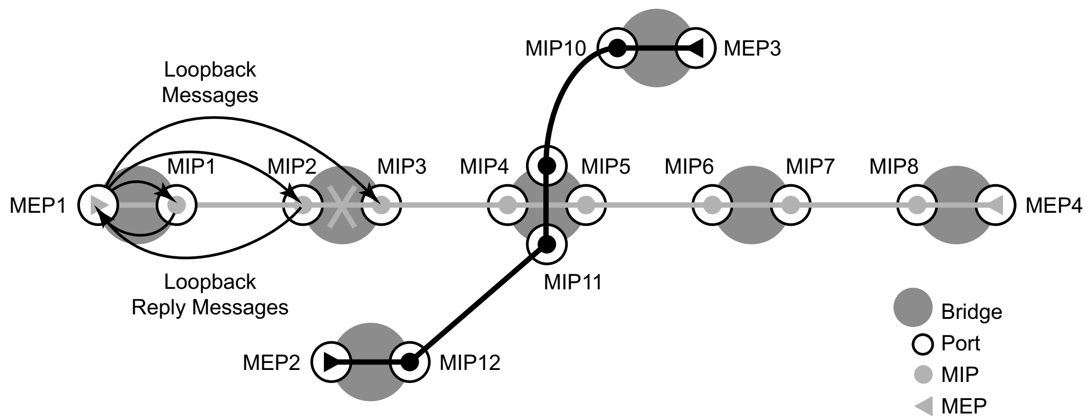

Figure 26: CFM Loopback

The following loopback-related functions are supported:

- Loopback message functionality on an MEP or MIP can be enabled or disabled.
- MEP — Supports generating loopback messages and responding to loopback messages with loopback reply messages.
- MIP — Supports responding to loopback messages with loopback reply messages when loopback messages are targeted to self.

- Displays the loopback test results on the originating MEP. There is a limit of ten outstanding tests per node.

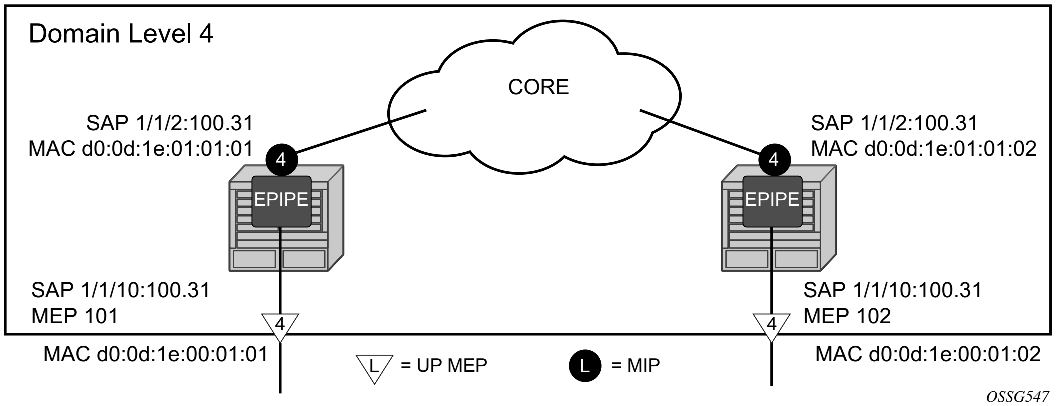

Figure 27: Loopback Configuration

```
# oam eth-cfm loopback d0:0d:1e:01:01:02 mep 101 domain 4 association
Eth-Cfm Loopback Test Initiated: Mac-Address: d0:0d:1e:01:01:02, out sap: 1/1/10:100.31
Sent 5 packets, received 5 packets [0 out-of-order, 0 Bad Msdu]

# oam eth-cfm loopback d0:0d:1e:00:01:02 mep 101 domain 4 association
Eth-Cfm Loopback Test Initiated: Mac-Address: d0:0d:1e:00:01:02, out sap: 1/1/10:100.31
Sent 5 packets, received 5 packets [0 out-of-order, 0 Bad Msdu]
```

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 27). Its function is similar to IP traceroute. Traces a specific MAC address through the service. The peer MEP responds with a linktrace reply message after successful inspection of the linktrace message. The MIPs along the path also process the linktrace message and respond with linktrace replies to the originating MEP if the received linktrace message that has a TTL greater than 1 and forward the linktrace message if a look up of the target MAC address in the Layer 2 FIB is successful. The originating MEP shall expect to receive multiple linktrace replies and from processing the linktrace replies, it can put together the route to the target bridge.

A traced MAC address is carried in the payload of the linktrace message, the target MAC. Each MIP and MEP receiving the linktrace message checks whether it has learned the target MAC address. In order to use linktrace the target MAC address must have been learned by the nodes in the network. If so, a linktrace message is sent back to the originating MEP. Also, a MIP forwards the linktrace message out of the port where the target MAC address was learned.

The linktrace message itself has a multicast destination address. On a broadcast LAN, it can be received by multiple nodes connected to that LAN. But, at most, one node will send a reply.

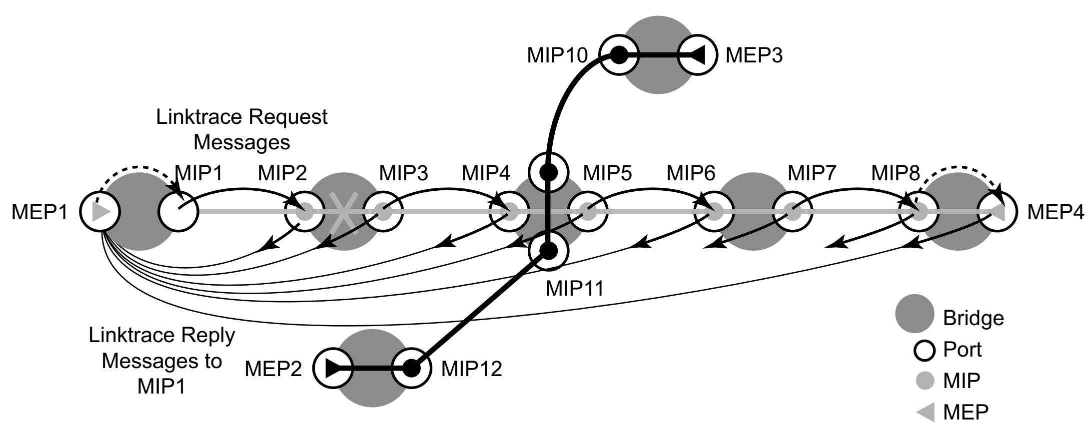

Figure 28: CFM Linktrace

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 when responding to the LTM. In mixed 802.ag and

Y.1731 environments the operator may chose to configure a Y.1731 context with an IEEE domain format.

The following linktrace related functions are supported:

- Enable or disables linktrace functions on an MEP.
- MEP — Supports generating linktrace messages and responding with linktrace reply messages.
- MIP — Supports responding to linktrace messages with linktrace reply messages when encoded TTL is greater than 1, and forward the linktrace messages accordingly if a lookup of the target MAC address in the Layer 2 FIB is successful.
- Displays linktrace test results on the originating MEP. There is a limit of ten outstanding tests per node. Storage is provided for up to ten MEPs and for the last ten responses. If more than ten responses are received older entries will be overwritten.

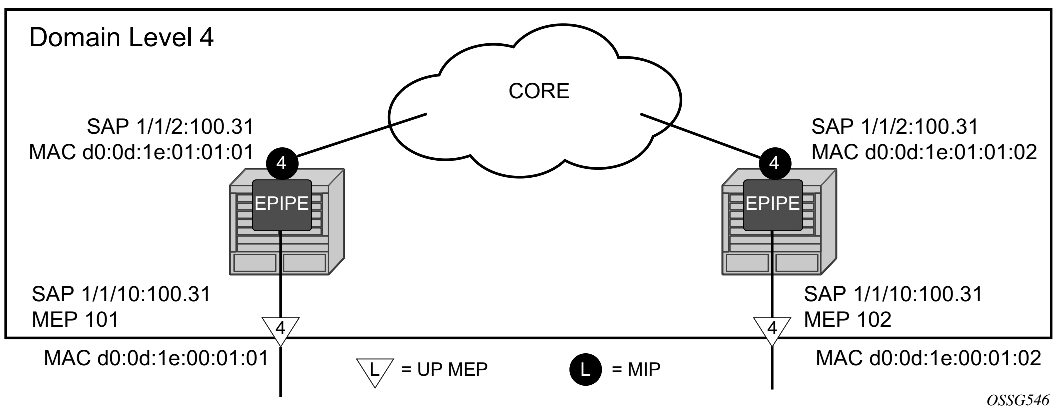

Figure 29: Linktrace Configuration

```
# oam eth-cfm linktrace d0:0d:1e:01:01:02 mep 101 domain 4 association 1

Index Ingress Mac          Egress Mac           Relay      Action
----- -------------------- -------------------- ---------- ----------
1     00:00:00:00:00:00    D0:0D:1E:01:01:01    n/a        forward
2     D0:0D:1E:01:01:02    00:00:00:00:00:00    n/a        none
----- -------------------- -------------------- ---------- ----------
No more responses received in the last 6 seconds.


# oam eth-cfm linktrace d0:0d:1e:00:01:02 mep 101 domain 4 association 1
```

```
Index Ingress Mac          Egress Mac           Relay      Action
----- -------------------- -------------------- ---------- ----------
1     00:00:00:00:00:00    D0:0D:1E:01:01:01    n/a        forward
2     D0:0D:1E:01:01:02    D0:0D:1E:00:01:02    n/a        terminate
----- -------------------- -------------------- ---------- ----------
No more responses received in the last 6 seconds.
```

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.

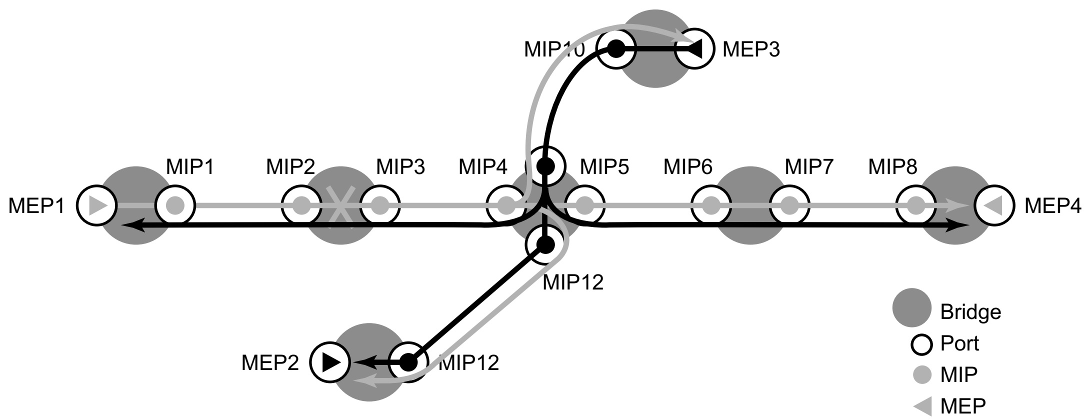

Figure 30: CFM Continuity Check

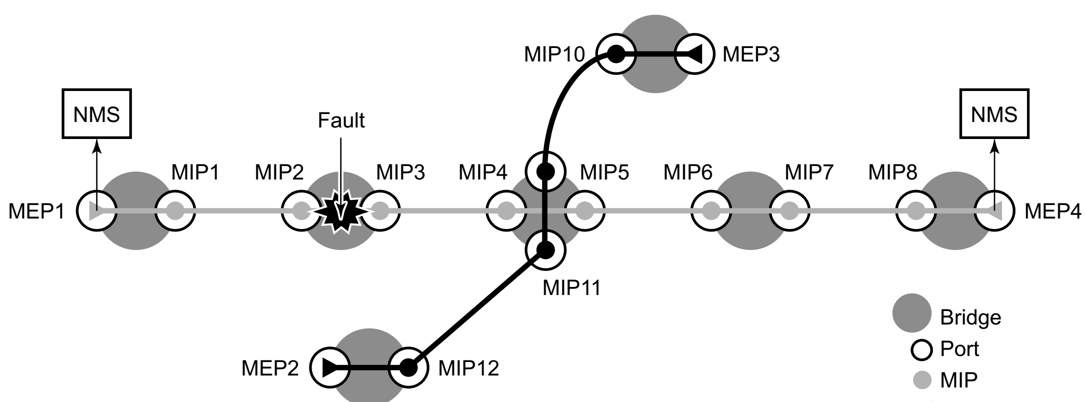

Figure 31: CFM CC Failure Scenario

The following functions are supported:

- Enable and disable CC for an MEP
- Configure and delete the MEP entries in the CC MEP monitoring database manually. It is only required to provision remote MEPs. Local MEPs shall be automatically put into the database when they are created.
- CCM transmit interval: 10ms, 100ms, 1s, 10s 60s, 600s. Default: 10s. Sub-second or fast CC requires a ESS-7/ESS-12 and SR-7/SR-12 with a minimum SF/CPM-3, with only a limited number supported on SF/CPM-1 & SF/CPM-2. When configuring MEPs with sub-second CCM intervals bandwidth consumption must be taken into consideration. Each CCM PDU is 100 bytes (800 bits). Taken individually this is a small value. However, the bandwidth consumption increases rapidly as multiple MEPs are configured with 10ms timers, 100 packets per second. Sub-second enabled MEPs are supported on the following:
  - → Down MEPs configured on Ethernet SAPs.
  - → Lowest MD-level, when multiple MEPs exist on same Ethernet SAP.
  - → Individual Ethernet tunnel paths requiring EAPs but not on the Ethernet tunnel itself. This requires a the MEPs to be part of the Y.1731 context because of the EAPS.
- CCM will declare a fault, when:
  - → The CCM stops hearing from one of the remote MEPs for 3.5 times CC interval
  - → Hears from a MEP with a LOWER MD level
  - → Hears from a MEP that is not part of the local MEPs MA
  - → Hears from a MEP that is in the same MA but not in the configured MEP list
  - → Hears from a MEP in the same MA with the same MEP id as the receiving MEP
  - → The CC interval of the remote MEP does not match the local configured CC interval
  - → The remote MEP is declaring a fault
- An alarm is raised and a trap is sent if the defect is greater than or equal to the configured low-priority-defect value.
- Remote Defect Indication (RDI) is supported but by default is not recognized as a defect condition because the low-priority-defect setting default does not include RDI.

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
            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. The following illustrates the output from a few of these display commands, taken from NODE1.

No defect conditions are raised. The **Defect** column in the first display is clear and the **Defect Flags** is the second display is also clear.

```
show eth-cfm cfm-stack-table
===============================================================================
CFM Stack Table Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM, A = AisRx

===============================================================================
CFM SAP Stack Table
===============================================================================
Sap                Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
1/1/2:100.31         4 Both         4          2  MIP d0:0d:1e:01:01:01 ------
1/1/10:100.31        4   Up         4          1  101 d0:0d:1e:00:01:01 ------
===============================================================================

show eth-cfm mep 101 domain 4 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index           : 4                          Direction         : Up
Ma-index           : 1                          Admin             : Enabled
MepId              : 101                        CCM-Enable        : Enabled
IfIndex            : 35979264                   PrimaryVid        : 2031716
Description        : (Not Specified)
FngState           : fngReset                   ControlMep        : False
LowestDefectPri    : macRemErrXcon              HighestDefect     : none
Defect Flags       : None
Mac Address        : d0:0d:1e:00:01:01          ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 1639                       CcmSequenceErr    : 0
Fault Propagation  : disabled                   FacilityFault     : n/a
MA-CcmInterval     : 1                          MA-CcmHoldTime    : 0ms
```

```
Eth-1Dm Threshold  : 3(sec)                  MD-Level          : 4
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled

Redundancy:
    MC-LAG State   : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================
```

The **all-remote-mepids** is the appropriate command to show the details for each configured peer, including the MAC address.

```
show eth-cfm mep 101 domain 4 association 1 all-remote-mepids

============================================================================
Eth-CFM Remote-Mep Table
============================================================================
R-mepId Rx CC  Rx Rdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
----------------------------------------------------------------------------
102     True   False  Up       Up     d0:0d:1e:00:01:02 02/02/2011 13:37:42
============================================================================
```

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.

Alarm Indication Signal (ETH-AIS Y.1731)

Alarm Indication Signal (AIS) provides an Y.1731 capable MEP the ability to signal a fault condition in the reverse direction of the MEP, out the passive side. When a fault condition is detected the MEP will generate AIS packets at the configured client levels and at the specified AIS interval until the condition is cleared. Currently a MEP configured to generate AIS must do so at a level higher than its own. The MEP configured on the service receiving the AIS packets is required to have the active side facing the receipt of the AIS packet and must be at the same level the AIS, The absence of an AIS packet for 3.5 times the AIS interval set by the sending node will clear the condition on the receiving MEP.

AIS generation is also not subject to the low-priority-defect setting. AIS, when enabled, generates when the MEP enters any defect condition, including RDI.

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, ETHAIS 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 ETHAIS information does not contain that information. Therefore, upon reception of a frame with ETH-AIS information, the MEP will suppress alarms for all peer MEPs whether there is still connectivity or not.

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 ETHAIS 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.

Specific configuration information required by a MEP to support ETH-AIS is the following:

- Client MEG Level — MEG level at which the most immediate client layer MIPs and MEPs exist.
- ETH-AIS transmission period — Determines transmission periodicity 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.

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 PW redundancy.

AIS is enabled under the service and has two parts, receive and transmit. Both of the 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. AIS is independent of the **low-priority-defect** setting, so that any fault in the MEP causes AIS to be generated.

```
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 and even though the **LowestDefectPri** setting is higher than the existing defect AIS is being transmitted. The **Eth-Ais Tx Counted** value is increasing, indicating that AIS is actively being sent.

```
# show eth-cfm mep 101 domain 4 association 1
===============================================================================
Eth-Cfm MEP Configuration Information
===============================================================================
Md-index          : 4                         Direction         : Up
Ma-index          : 1                         Admin             : Enabled
MepId             : 101                       CCM-Enable        : Disabled
IfIndex           : 35979264                  PrimaryVid        : 2031716
Description       : (Not Specified)
FngState          : fngReset                  ControlMep        : False
LowestDefectPri   : macRemErrXcon             HighestDefect     : none
Defect Flags      : bDefRDICCM
Mac Address       : d0:0d:1e:00:01:01         ControlMep        : False
CcmLtmPriority    : 7
CcmTx             : 2578                      CcmSequenceErr    : 0
Fault Propagation  : disabled                 FacilityFault     : n/a
MA-CcmInterval    : 1                         MA-CcmHoldTime    : 0ms
Eth-1Dm Threshold  : 3(sec)                   MD-Level          : 4
Eth-Ais:          : Enabled                   Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 288
Eth-Ais Tx Levels  : 5
Eth-Tst:          : Disabled

Redundancy:
    MC-LAG State   : n/a

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================
```

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
===============================================================
                                 Current        Accumulate
                  FrameCount     ErrBits        ErrBits
Peer Mac Addr     ByteCount      CrcErrs        CrcErrs
---------------------------------------------------------------
d0:0d:1e:00:01:01 3              0              0
                  3000           0              0
===============================================================
```

One-Way Delay Measurement (ETH-1DM Y.1731)

One-way delay measurement allows the operator the ability to check unidirectional delay between MEPs. 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 and SR-1, will display different results than hardware time stamp capable devices, like the SR-7/SR-12 and ESS-7/ESS-12.

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
===================================================================
```

Synthetic Loss Measurement (ETH-SL)

Alcatel-Lucent applied pre-standard OpCodes 53 (Synthetic Loss Reply) and 54 (Synthetic Loss Message) for the purpose of measuring loss using synthetic packets.

Notes: These will be changes to the assigned standard values in a future release. This means that the Release 9.0R1 is pre-standard and will not interoperate with future releases of SLM/SLR that support the standard OpCode values.

This synthetic loss measurement approach is a single-ended feature that allows the operator to run on-demand and proactive tests to determine “in”, “out” loss and “unacknowledged” packets. This approach can be used between peer MEPs in both point to point and multipoint services. Only remote MEP peers within the association and matching the unicast destination will respond to the SLM packet.

The specification uses various sequence numbers in order to determine in which direction the loss occurred. 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. Facility based MEPs, although not blocked, do not support the ETH-SL functions. Executing these features on Facility MEPs may provide inconsistent measurement. Support for Facility MEPs is being addressed in a near term release.

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 is an acceptable approach.

There is no way for the responding node to understand when a test is completed. For this reason a configurable “inactivity-timer” determines the length of time a test is valid. The timer will maintain an active test as long as it is receiving packets for that specific test, defined by the test-id, remote MEP Id and source MAC. When there is a gap between the packets that exceeds the inactivity-timer the responding node will respond with a sequence number of one regardless of what the sequence number was the instantiating node sent. This means the remote MEP believes the previous test has expired and these probes are part of a new test. The default for the inactivity-timer is 100 second and has a range of ten to 100 seconds.

The responding node will be limited to 1000 concurrent test SLM tests. Any test that attempts to involve a node that is already actively processing 1000 SLM tests will show up as “out loss” or “unacknowledged” packets on the node that instantiated the test because the packets will be silently discarded at the responder. It is important for the operator to understand this is silent and no log entries or alarms will be raised. It is also important to keep in mind that these packets are ETH-CFM based and the different platforms stated receive rate for ETH-CFM must not be exceeded.

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 proactively trend packet loss for service based MEPs.

Configuration Example

The following illustration, , shows the configuration required for proactive SLM test using SAA.

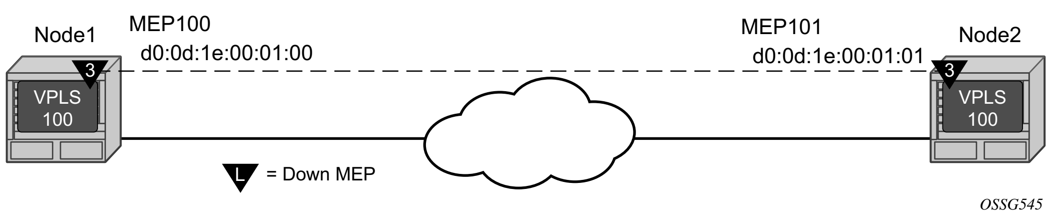

Figure 32: SLM Example

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
===============================================================================
```

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 or ATM.

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.

CFM Connectivity Fault Conditions

CFM MEP declares a connectivity fault when its defect flag is equal to or higher than its configured lowest defect priority. The defect can be any of the following depending on configuration:

- DefRDICCM
- DefMACstatus
- DefRemoteCCM
- DefErrorCCM
- DefXconCCM

The following additional fault condition applies to Y.1731 MEPs:

- Reception of AIS for the local MEP level

Setting the lowest defect priority to allDef may cause problems when fault propagation is enabled in the MEP. In this scenario, when MEP A sends CCM to MEP B with interface status down, MEP B will respond with a CCM with RDI set. If MEP A is configured to accept RDI as a fault, then it gets into a dead lock state, where both MEPs will declare fault and never be able to recover. The default lowest defect priority is DefMACstatus, which will not be a problem when interface status

TLV is used. It is also very important that different Ethernet OAM strategies should not overlap the span of each other. In some cases, independent functions attempting to perform their normal fault handling can negatively impact the other. This interaction can lead to fault propagation in the direction toward the original fault, a false positive, or worse, a deadlock condition that may require the operator to modify the configuration to escape the condition. For example, overlapping Link Loss Forwarding (LLF) and ETH-CFM fault propagation could cause these issues.

For the DefRemoteCCM fault, it is raised when any remote MEP is down. So whenever a remote MEP fails and fault propagation is enabled, a fault is propagated to SMGR.

CFM Fault Propagation Methods

When CFM is the OAM module at the other end, it is required to use any of the following methods (depending on local configuration) to notify the remote peer:

- Generating AIS for certain MEP levels
- Sending CCM with interface status TLV “down”
- Stopping CCM transmission

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 must be done instantly without waiting for the next CCM transmit interval. This rule applies to CFM fault notification for all services.

Notifications from SMGR to the CFM MEPs for fault propagation should include a direction for the propagation (up or down: up means in the direction of coming into the SAP/SDP-binding; down means in the direction of going out of the SAP/SDP-binding), so that the MEP knows what method to use. For instance, an up fault propagation notification to a down MEP will trigger an AIS, while a down fault propagation to the same MEP can trigger a CCM with interface TLV with status down.

For a specific SAP/SDP-binding, CFM and SMGR can only propagate one single fault to each other for each direction (up or down).

When there are multiple MEPs (at different levels) on a single SAP/SDP-binding, the fault reported from CFM to SMGR will be the logical OR of results from all MEPs. Basically, the first fault from any MEP will be reported, and the fault will not be cleared as long as there is a fault in any local MEP on the SAP/SDP-binding.

Epipe Services

Down and up MEPs are supported for Epipe services as well as fault propagation. When there are both up and down MEPs configured in the same SAP/SDP-binding and both MEPs have fault propagation enabled, a fault detected by one of them will be propagated to the other, which in turn will propagate fault in its own direction.

CFM Detected Fault

When a MEP detects a fault and fault propagation is enabled for the MEP, CFM needs to communicate the fault to SMGR, so SMGR will mark the SAP/SDP-binding faulty but still oper-up. CFM traffic can still be transmitted to or received from the SAP/SDP-binding to ensure when the fault is cleared, the SAP will go back to normal operational state. Since the operational status of the SAP/SDP-binding is not affected by the fault, no fault handling is performed. For example, applications relying on the operational status are not affected.

If the MEP is an up MEP, the fault is propagated to the OAM components on the same SAP/SDP-binding; if the MEP is a down MEP, the fault is propagated to the OAM components on the mate SAP/SDP-binding at the other side of the service.

SAP/SDP-Binding Failure (Including Pseudowire Status)

When a SAP/SDP-binding becomes faulty (oper-down, admin-down, or pseudowire status faulty), SMGR needs to propagate the fault to up MEP(s) on the same SAP/SDP-bindings about the fault, as well as to OAM components (such as down MEPs and E-LMI) on the mate SAP/SDP-binding.

Service Down

This section describes procedures for the scenario where an Epipe service is down due to the following:

- Service is administratively shutdown. When service is administratively shutdown, the fault is propagated to the SAP/SDP-bindings in the service.
- If the Epipe service is used as a PBB tunnel into a B-VPLS, the Epipe service is also considered operationally down when the B-VPLS service is administratively shutdown or operationally down. If this is the case, fault is propagated to the Epipe SAP.
- In addition, one or more SAPs/SDP-bindings in the B-VPLS can be configured to propagate fault to this Epipe (see fault-propagation-bmac below). If the B-VPLS is operationally up but all of these entities have detected fault or are down, the fault is propagated to this Epipe’s SAP.

Interaction with Pseudowire Redundancy

When a fault occurs on the SAP side, the pseudowire status bit is set for both active and standby pseudowires. When only one of the pseudowire is faulty, SMGR does not notify CFM. The notification occurs only when both pseudowire becomes faulty. The SMGR propagates the fault to CFM.

Since there is no fault handling in the pipe service, any CFM fault detected on an SDP binding is not used in the pseudowire redundancy's algorithm to choose the most suitable SDP binding to transmit on.

Ipipe Services

For Ipipe services, only down MEPs are supported on Ethernet SAPs.

CFM Detected Fault

When a MEP detects a fault and fault propagation is enabled for the MEP, CFM needs to communicate the fault to SMGR, so SMGR will mark the SAP/SDP-binding faulty but still oper-up. CFM traffic can still be transmitted to or received from the SAP/SDP-binding to ensure when the fault is cleared, the SAP will go back to normal operational state.

Because the MEP is a down MEP, the fault is always propagated to the OAM components on the mate SAP/SDP-binding at the other side of the service.

SAP/SDP-binding Failure (Including Pseudowire Status)

When a SAP/SDP-binding becomes faulty (oper-down, admin-down, or pseudowire status faulty), SMGR propagates the fault to OAM components on the mate SAP/SDP-binding.

Service Administratively Shutdown

When the service is administratively shutdown, SMGR propagates the fault to OAM components on both SAP/SDP-bindings.

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.

VPLS Service

For VPLS services, on down MEPs are supported for fault propagation.

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 on page 185, 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 m VPLS 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.

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.

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.

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.

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.

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.

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.

LLF and CFM Fault Propagation

LLF and CFM fault propagation are mutually exclusive. CLI protection is in place to prevent enabling both LLF and CFM fault propagation in the same service, on the same node and at the same time. However, there are still instances where irresolvable fault loops can occur when the two schemes are deployed within the same service on different nodes. This is not preventable by the CLI. At no time should these two fault propagation schemes be enabled within the same service.

802.3ah EFM OAM Mapping and Interaction with Service Manager

802.3ah EFM OAM declares a link fault when any of the following occurs:

- Loss of OAMPDU for a certain period of time
- Receiving OAMPDU with link fault flags from the peer

When 802.3ah EFM OAM declares a fault, the port goes into operation state down. The SMGR communicates the fault to CFM MEPs in the service.

OAM fault propagation in the opposite direction (SMGR to EFM OAM) is not supported.

Service Assurance Agent Overview

In the last few years, service delivery to customers has drastically changed. Services such as VPLS and VPRN are offered. The introduction of Broadband Service Termination Architecture (BSTA) applications such as Voice over IP (VoIP), TV delivery, video and high speed Internet services force carriers to produce services where the health and quality of Service Level Agreement (SLA) commitments are verifiable to the customer and internally within the carrier.

SAA is a feature that monitors network operations using statistics such as jitter, latency, response time, and packet loss. The information can be used to troubleshoot network problems, problem prevention, and network topology planning.

The results are saved in SNMP tables are queried by either the CLI or a management system. Threshold monitors allow for both rising and falling threshold events to alert the provider if SLA performance statistics deviate from the required parameters.

SAA Application

SAA allows two-way timing for several applications. This provides the carrier and their customers with data to verify that the SLA agreements are being properly enforced.

Two-way time measures requests from this node to the specified DNS server. This is done by performing an address request followed by an immediate release of the acquired address once the time measurement has been performed.

Traceroute Implementation

Various applications, such as lsp-trace, traceroute and vprn-trace, and traceroute, pass through the P-chip on the way to the control CPU. At this point, and when it egresses the control CPU, the P-chip should insert a timestamp inside the packet. Only packets processed by the Control CPU are processed.

When interpreting these timestamps care must be taken that some nodes are not capable of providing timestamps, as such timestamps must be associated with the same IP-address that is being returned to the originator to indicate what hop is being measured.

NTP

Because NTP precision can vary (+/- 1.5ms between nodes even under best case conditions), SAA one-way latency measurements might display negative values, especially when testing network segments with very low latencies. The one-way time measurement relies on the accuracy of NTP between the sending and responding nodes.

Ethernet CFM

Loopback (LBM), linktrace (LTR) and two-way-delay measurements (Y.1731 ETH-DMM) can be scheduled using SAA. Additional timestamping is required for non Y.1731 delay-measurement tests, to be specific, loopback and linktrace tests. An organization-specific TLV is used on both sender and receiver nodes to carry the timestamp information. Currently, timestamps are only applied by the sender node. This means any time measurements resulting from loopback and linktrace tests includes the packet processing time of the remote node. Since Y.1731 ETH-DMM uses a four time stamp approach to remove the remote processing time it should be used for accurate delay measurements.

The SAA versions of the CFM loopback, linktrace and ETH-DMM tests support send-count, interval, timeout, and FC. The existing CFM OAM commands have not been extended to support send-count and interval natively. The summary of the test results are stored in an accounting file that is specified in the SAA accounting-policy.

Writing SAA Results to Accounting Files

SAA statistics enables writing statistics to an accounting file. When results are calculated an accounting record is generated.

In order to write the SAA results to an accounting file in a compressed XML format at the termination of every test, the results must be collected, and, in addition to creating the entry in the appropriate MIB table for this SAA test, a record must be generated in the appropriate accounting file.

Accounting File Management

Because the SAA accounting files have a similar role to existing accounting files that are used for billing purposes, existing file management information is leveraged for these accounting (billing) files.

Assigning SAA to an Accounting File ID

Once an accounting file has been created, accounting information can be specified and will be collected by the **config>log>acct-policy> to file** *log-file-id* context.

Continuous Testing

When you configure a test, use the **config**>**saa**>**test**>**continuous** command to make the test run continuously. Use the **no continuous** command to disable continuous testing and **shutdown** to disable the test completely. Once you have configured a test as continuous, you cannot start or stop it by using the **saa** *test-name* [**owner** *test-owner*] {**start** | **stop**} [**no-accounting**] command.

Configuring SAA Test Parameters

The following example displays an SAA configuration:

```
A:ALA-48>config>saa# info
----------------------------------------------------------------
...
        test "vprnTr"
            type
                vprn-trace 5 source 20.20.12.1 destination 10.10.12.2
            exit
            jitter-event rising-threshold 5 falling-threshold 2 inbound
            jitter-event rising-threshold 5 falling-threshold 2 outbound
            jitter-event rising-threshold 6 falling-threshold 3
            loss-event rising-threshold 1 inbound
            loss-event rising-threshold 1 outbound
            loss-event rising-threshold 1
            latency-event rising-threshold 30 falling-threshold 1 inbound
            latency-event rising-threshold 30 falling-threshold 1 outbound
            latency-event rising-threshold 30 falling-threshold 1
            no shutdown
        exit
----------------------------------------------------------------
A:ALA-48>config>saa#
```

After running the test twice, the result is displayed below:

```
*A:ALA-48>config>saa# show saa vprnTr
===============================================================================
SAA Test Information
===============================================================================
Test name                    : vprnTr
Owner name                   : TiMOS CLI
Description                  : N/A
Accounting policy            : None
Administrative status        : Enabled
Test type                    : vprn-trace 5 source 20.20.12.1 destination 10.
                               10.12.2
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    5.00       None       Never               None
            Falling   2.00       None       Never               None
Jitter-out  Rising    5.00       None       Never               None
            Falling   2.00       None       Never               None
Jitter-rt   Rising    6.00       None       Never               None
            Falling   3.00       None       Never               None
Latency-in  Rising    30.0       None       Never               None
            Falling   1.00       None       Never               None
Latency-out Rising    30.0       None       Never               None
            Falling   1.00       None       Never               None
Latency-rt  Rising    30.0       None       Never               None
```

```
            Falling   1.00       None       Never               None
Loss-in     Rising    1          None       Never               None
            Falling   None       None       Never               None
Loss-out    Rising    1          None       Never               None
            Falling   None       None       Never               None
Loss-rt     Rising    1          None       Never               None
            Falling   None       None       Never               None
===============================================================================
*A:ALA-48>config>saa#
```

Configuring Trap Generation

The following shows an example of a ping test.

```
*A:bksim130>config>saa>test>trap-gen# probe-fail-enable
*A:bksim130>config>saa>test>trap-gen# probe-fail-threshold 3
*A:bksim130>config>saa>test>trap-gen# test-completion-enable
*A:bksim130>config>saa>test>trap-gen# test-fail-enable
*A:bksim130>config>saa>test>trap-gen# test-fail-threshold 2

*A:bksim130>config>saa>test>trap-gen# info
----------------------------------------------
            trap-gen
                probe-fail-enable
                probe-fail-threshold 3
                test-completion-enable
                test-fail-enable
                test-fail-threshold 2
            exit
----------------------------------------------
*A:bksim130>config>saa>test>trap-gen#
```

The following shows an example of a trap generation configuration.

```
*A:bksim130# configure saa test mySaaTraceRouteTest1
*A:bksim130>config>saa>test$ type icmp-trace 11.22.33.44

*A:bksim130>config>saa>test$ trap-gen
*A:bksim130>config>saa>test>trap-gen$ probe-fail-enable
MINOR: CLI SAA test with testName=mySaaTraceRouteTest1, ownerName=TiMOS CLI:  probe-fail-
enable applies to ping tests only.

*A:bksim130>config>saa>test>trap-gen$ probe-fail-threshold 2
MINOR: CLI SAA test with testName=mySaaTraceRouteTest1, ownerName=TiMOS CLI:  probe-fail-
threshold applies to ping tests only.

*A:bksim130>config>saa>test>trap-gen$ test-completion-enable
*A:bksim130>config>saa>test>trap-gen$ test-fail-enable
*A:bksim130>config>saa>test>trap-gen$ test-fail-threshold 2
MINOR: CLI SAA test with testName=mySaaTraceRouteTest1, ownerName=TiMOS CLI:  test-fail-
threshold applies to ping tests only.

*A:bksim130>config>saa>test>trap-gen$ info
----------------------------------------------
            trap-gen
                test-completion-enable
                test-fail-enable
            exit
----------------------------------------------
*A:bksim130>config>saa>test>trap-gen#
```

Diagnostics Command Reference

OAM Commands

Base Operational Commands

GLOBAL
- — **ping** [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address* | *ipv6-address* | *dns-name*] [**interval** *seconds*] [{**next-hop** *ip-address*} | {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance*] [**timeout** *timeout*]
- — **traceroute** [*ip-address* | *dns-name*] [**ttl** *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**]

ATM Diagnostics

GLOBAL
- — **oam**
  - — **atm-ping** *port-id:vpi/vci* [**end-to-end** | **segment**] [**dest** *destination-id*][**send-count** *send-count*][**timeout** *seconds*][**interval** *seconds*]

IGMP Snooping

GLOBAL
- — **oam**
  - — **mfib-ping** **service** *service-id* **source** *src-ip* **destination** *mcast-address* [**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**send-count** *send-count*] [**timeout** *timeout*]

LDP Diagnostics

GLOBAL
- — **oam**
  - — **ldp-treetrace** {**prefix** *ip-prefix/mask*} [**max-ttl** *ttl-value*] [**max-path** *max-paths*] [**timeout** *tim-eout*] [**retry-count** *retry-count*] [**fc** *fc-name* [**profile** *profile*]]
- — **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-time-stamp-format** {**rfc4379** | **unix**}

TWAMP

GLOBAL
— **oam**
  — **oam-test**
    — **twamp**
      — **server**
        — [**no**] **prefix** {*address/prefix-length* | *address netmask*}
          —[**no**] **description** *description*
          —[**no**] **max-conn-prefix** *count*
          —[**no**] **max-sess-prefix** *count*
          —**no**] **shutdown**
        — [**no**] **inactivity-timeout** *seconds*
        — [**no**] **max-conn-server** *count*
        — [**no**] **max-sess-server** *count*
        — [**no**] **port** *number*
        — [**no**] **shutdown**

LSP Diagnostics

GLOBAL
— **oam**
  — **lsp-ping** {{[*lsp-name*] [**path** *path-name*]} | {**prefix** ip-prefix/mask}} [**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**]
  — **lsp-trace** {{[*lsp-name]* [**path** path-name]} | {**prefix** ip-prefix/mask}} [**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**]

— **p2mp-lsp-ping** *lsp-name* [**p2mp-instance** *instance-name* [**s2l-dest-address** *ip-address* [...(up to 5 max)]]] [**fc** *fc-name* [**profile** {**in**|**out**}]] [**size** *octets*] [**ttl** label-ttl] [**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**]

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*]

Common Service Diagnostics

GLOBAL

— **oam**

- — **ancp** {**subscriber** *sub-ident-string* | **ancp-string** *ancp-string*} **loopback** [**count** *count*] [**time-out** *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*] [**tim-eout** *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*]

VLL Diagnostics

GLOBAL

— **oam**

- — **vccv-ping** *sdp-id:vc-id* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*][**reply-mode** {**ip-routed** | **control-channel**}][**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*][**ttl** *vc-label-ttl*]
- — **vccv-trace** *sdp-id:vc-id* [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**reply-mode** *ip-routed|countrol-channel*] [**probe-count** *probes-per-hop*] [**timeout** *timeout*] [**interval** *interval*] [**min-ttl** *min-vc-label-ttl*] [**max-ttl** max-*vc-label-ttl*] [**max-fai**l *no-response-count*] [**detail**]

VPLS MAC Diagnostics

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*] [**send-control**] [**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*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]
- — **mac-populate** *service-id* **mac** *ieee-address* [**flood**] [**age** *seconds*] [**force**] [**target-sap** *sap-id*] [**send-control**]
- — **mac-purge** *service-id* **target** *ieee-address* [**flood**] [**send-control**] [**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*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]

Ethernet in the First Mile (EFM) Commands

GLOBAL

— **oam**

- — **efm** *port-id* **local-loopback** {**start** | **stop**}
- — **efm** *port-id* **remote-loopback** {**start** | **stop**}

ETH-CFM OAM Commands

oam

- — **eth-cfm eth-test** ***mac-address*** **mep** ***mep-id*** **domain** ***md-index*** **association** ***ma-index*** [**priority** *priority*] [**data-length** *data-length*]
- — **eth-cfm linktrace** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**ttl** *ttl-value*]
- — **eth-cfm loopback** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**send-count** *send-count*] [**size** *data-size*] [**priority** *priority*]
- — **eth-cfm one-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]
- — **eth-cfm two-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]
- — **eth-cfm two-way-slm-test** ***mac-address*** **mep** ***mep-id*** **domain** ***md-index*** **association** ***ma-index*** **[priority** *priority*]] **[send-count** ***send-count*** **][size data-size][timeout** ***timeout*****] [interval** ***interval*****]**

SAA Commands

config
- — **saa**
  - — [**no**] **test** *test-name* [**owner** *test-owner*]
    - — **accounting-policy** *acct-policy-id*
    - — **no accounting-policy**
    - — [**no**] **continuous**
    - — **description** *description-string*
    - — **no description**
    - — [**no**] **jitter-event** **rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
    - — [**no**] **latency-event** **rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
    - — [**no**] **loss-event** **rising-threshold** *threshold* [**falling-threshold** *threshold*] [*direction*]
    - — [**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**] [**interval** *interval*]
      - — **dns** **target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]
      - — **eth-cfm-linktrace** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**ttl** *ttl-value*] [**fc** {*fc-name*} [**profile** {**in**|**out**}]] [**send-count** send-count] [**timeout** *timeout*] [**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*][**timeout** *timeout*] [**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*][**timeout** *timeout*] [**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][**timeout** *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*] [**timeout** *timeout*] [**interval** *interval*]
      - — **icmp-ping** [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address* | *dns-name*] [**interval** *seconds*] [{**next-hop** *ip-address*}|{**interface** *interface-name*}|**bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance*] [**timeout** *timeout*]
      - — **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*]}|{**prefix** *ip-prefix/mask*}} [**src-ip-address** *ip-addr*] [**size** *octets*] [**ttl** *label-ttl*] [**timeout** *timeout*] [**interval** *interval*] [**fc** *fc-name* [**profile** {in | out}]] [**send-count** *send-count*]
— **lsp-trace** {{*lsp-name* [**path** *path-name*]}|{**prefix** *ip-prefix/mask*} }[**src-ip-address** *ip-addr*] [**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*]]
— **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*] [**timeout** *timeout*]
— **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*] [**timeout** *timeout*]
— **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**send-count** *send-count*][**timeout** *seconds*] [**interval** *seconds*]
— **vccv-ping** *sdp-id:vc-id* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*][**reply-mode** {**ip-routed** | **control-channel**}][**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**send-count** *send-count*][**timeout** *timeout*] [**interval** *interval*][**ttl** *vc-label-ttl*]
— **vccv-trace** *sdp-id:vc-id* [**size** *octets*][**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*][**max-fail** *no-response-count*][**probe-count** *probe-count*][**reply-mode** *ip-routed*|**control-channel**][**timeout** *timeout-value*][**interval** *interval-value*][**fc** *fc-name* [**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**] [**timeout** *timeout*] [**interval** *seconds*]
— **vprn-trace** *service-id* **source** *src-ip* **destination** *dst-ip* [**fc** *fc-name* [**profile** **in** | **out**]] [**size** *size*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**probe-count** *send-count*] [**return-control**] [**timeout** *timeout*] [**interval** *interval*]

Show Commands

show
- — **eth-cfm**
  - — **association** [*ma-index*] [**detail**]
  - — **cfm-stack-table** [**port** [*port-id* [**vlan** *vlan-id*]]| **sdp** *sdp-id*[:*vc-id*]][**level** 0..7] [**direction up** | **down**]
  - — **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**[.*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 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**] [**detail**]
  - — **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]
  - — **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**remote-mepid** *mep-id* | **all-remote-mepids**]
  - — **mep** *mep-id* **domain** *md-index* **association** *ma-index* **eth-test-results** [**remote-peer** *mac-address*]
  - — **mep** *mep-id* **domain** *md-index* **association** *ma-index* **one-way-delay-test** [**remote-peer** *mac-address*]
  - — **mep** *mep-id* **domain** *md-index* **association** *ma-index* **two-way-delay-test** [**remote-peer** *mac-address*]
  - — **mep** *mep-id* **domain** *md-index* **association** *ma-index* **two-way-slm-test** [**remote-peer** *mac-address*]
  - — **mip**
  - — **system-config**
- — **saa** [*test-name* [**owner** *test-owner*]]
- — **test-oam**
  - — **ldp-treetrace** [**prefix** *ip-prefix/mask*] [**detail**]

Clear Commands

clear
- — **saa** [*test-name* [**owner** *test-owner*]]

Debug Commands

debug
- — **oam**

— **lsp-ping-trace** [**tx** | **rx** | **both**] [**raw** | **detail**]
— **no lsp-ping-trace**

OAM and SAA Commands

Command Hierarchies

Operational 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.

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 ipv4-a-record is specified, dns-names are queried for A-records only. If ipv6-aaaa-record is specified, AAAA-records are queried first, and if a successful reply is not received, the dns-server is queried for A-records.

Parameters

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **send-count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

ip-address — The IP or IPv6 address of the primary DNS server.

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..FFFF]H

d - [0..255]D

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 120

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

record-type **—** Specifies a record type.

Values **ipv4-a-record —** A record specific mapping a host name to an IPv4 address.
ipv6-aaaa-record — A record specific to the Internet class that stores a single IPv6 address.

ping

Syntax **ping** [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address* | *dns-name*] [**interval** *seconds*] [{**next-hop** *ip-address*} |

{**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance*] [**timeout** *timeout*]

Context <GLOBAL>

Description This command verifies the reachability of a remote host.

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[-interface]
x:x:x:x:x:x:d.d.d.d[-interface]
x: [0 — FFFF]H
d: [0 — 255]D
interface:32 characters maximum, mandatory for link local addresses

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 — 128

tos *type-of-service* **—** Specifies the service type.

Values 0 — 255

size *bytes* **—** The request packet size in bytes, expressed as a decimal integer.

Values 0 — 16384

pattern *pattern* **—** The date portion in a ping packet will be filled with the pattern value specified. If not specified, position info will be filled instead.

Values 0 — 65535

source *ip-address* **—** Specifies the IP address to be used.

Values
ipv4-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x:x
x:x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

router *router-instance* **—** Specifies the router name or service ID.

Values
router-name: Base , management
service-id: 1 — 2147483647

Default Base

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:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]

count *requests* **—** Specifies the number of times to perform an OAM ping probe operation. Each OAM echo message request must either timeout or receive a reply before the next message request is sent.

Values 1 — 100000

Default 5

do-not-fragment — Sets the DF (Do Not Fragment) bit in the ICMP ping packet.

timeout *seconds* **—** Overrides the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

traceroute

Syntax **traceroute** [*ip-address* | *dns-name*] [**ttl** *ttl*] [**wait** *milli-seconds*] [**no-dns**] [**source** *ip-address*] [tos *type-of-service*] [**router** *router-instance*]

Context oam

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:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

d*ns-name* — The DNS name of the far-end device to which to send the traceroute request message, expressed as a character string.

ttl *ttl* **—** The maximum Time-To-Live (TTL) value to include in the traceroute request, expressed as a decimal integer.

Values 1 — 255

wait *milliseconds* **—** The time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Default 5000

Values 1 — 60000

no-dns — When the **no-dns** keyword is specified, DNS lookups of the responding hosts will not be performed, only the IP addresses will be printed.

Default DNS lookups are performed

source *ip-address* **—** The source IP address to use as the source of the probe packets in dotted decimal notation. If the IP address is not one of the device's interfaces, an error is returned.

tos *type-of-service* **—** The type-of-service (TOS) bits in the IP header of the probe packets, expressed as a decimal integer.

Values 0 — 255

router *router-name* **—** Specify the alphanumeric character string up to 32 characters.

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 on which this service is defined.

Values 1 — 2147483647

p2mp-lsp-ping

Syntax **p2mp-lsp-ping** {{*lsp-name* [**p2mp-instance** *instance-name* [**s2l-dest-address** *ip-address* [...(up to 5 max)]]]} | {**ldp** *p2mp-identifier* [**sender-addr** *ip-address*] [**leaf-addr** *ip-address*…[ip-address…*up to 5 max*]]}} [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**ttl** *label-ttl*] [**timeout** *timeout*] [**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 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.

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 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 router that receives the message request. The egress mappings of the egress network interface on the far-end far-end 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 far-end.

Default be

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

ldp ***p2mp-identifier*** **—** Identifier to specify a LDP P2MP LSP to ping.

Values The p2mp-identifier must be a 32 bit integer.

leaf-addr *ip-address…*[*ip-address…up to 5max*] **—** Specifies the list of egress LER system addresses which are required to reply to LSP ping echo request message.

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.

size *octets* **—** The MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 128 octets. The system sends the minimum packet size for an RSVP P2MP LSP.

Values 128 — 65535

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 — 120

ttl *label-ttl* **—** The TTL value for the MPLS label, expressed as a decimal integer.

Default 255

Values 1 — 255

p2mp-lsp-trace

Syntax **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**]

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 braches 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 B-flag in the downstream mapping TLV to indicate to the sender of the echo request message it has other branches for this LSP. A bud LSR will also set the E-flag in the downstream mapping TLV to indicate to the sender of the echo request message that it is also an egress LER for the P2MP LSP when the traced egress is reachable via a downstream branch. In this case, the return code must correspond to the LSR role and must code #8: “Label switched at stack-depth <RSC>”.

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** 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 note that receives the message request. The egress mappings of the egress network interface on the far-end node 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 node.

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 — 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 — 255

max-ttl *max-label-ttl* — the maximum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 30

Values 1-255

min-ttl *min-label-ttl* — The minimum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

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* **—** The number of LSP trace echo request messages to send per TTL value.

Default 1

Values 1 — 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 MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 128 octets. The system sends the minimum packet size for an RSVP P2MP LSP.

Values 128 — 65535

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 — 60

ATM Diagnostics

atm-ping

Syntax **atm-ping** *port-id*: *vpi/vci* [**end-to-end | segment**] [**dest** *destination-id*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *seconds*]

Context <GLOBAL>

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

port-id	*slot/mda/port*	
aps-id	aps-*group-id*	
	aps	keyword
	group-id	1 — 64
vpi	0 — 4095 (NNI)	
	0 — 255 (UNI)	
vci	1, 2, 5 — 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 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is

used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

Service Diagnostics

ancp

Syntax **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**]

Context <GLOBAL>

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** ***sub-ident-string*** **—** Specifies an existing subscriber-id. The node will use the ancp-string associated with the provided subscriber-id to identify the circuit.

ancp-string *ancp-string* **—** 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.

1 — 32

timeout *seconds* **—** Specifies how long the controlling node will wait for a result.

0 — 300

alarm — Specifies that the CLI the result will be retuned 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 7750 SR. 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 7750 SR.

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 — 17407

size-inc *start-octets end-octets* **—** Indicates an incremental path MTU test will be performed with by sending a series of message requests with increasing MTU sizes. The *start-octets* and *end-octets* parameters are described below.

start-octets **—** The beginning size in octets of the first message sent for an incremental MTU test, expressed as a decimal integer.

Values 40 — 9198

end-octets — The ending size in octets of the last message sent for an incremental MTU test, expressed as a decimal integer. The specified value must be greater than *start-octets*.

Values 40 — 9198

step *step-size* **—** The number of octets to increment the message size request for each message sent for an incremental MTU test, expressed as a decimal integer. The next size message will not be sent until a reply is received or three messages have timed out at the current size.

If the incremented size exceeds the *end-octets* value, no more messages will be sent.

Default 32

Values 1 — 512

timeout *seconds* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *seconds* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

Output **Sample SDP MTU Path Test Sample Output**

```
*A:Dut-A# oam sdp-mtu 1201 size-inc 512 3072 step 256
Size    Sent    Response
----------------------------
512     .        Success
768     .        Success
1024    .        Success
1280    .        Success
1536    .        Success
1792    .        Success
2048    .        Success
2304    .        Success
2560    .        Success
2816    .        Success
3072    .        Success

Maximum Response Size: 3072
*A:Dut-A#
```

svc-ping

Syntax **svc-ping** *ip-address* [**service** *service-id*] [**local-sdp**] [**remote-sdp**]

Context <GLOBAL>

Description Tests a service ID for correct and consistent provisioning between two service end points.

The **svc-ping** command accepts a far-end IP address and a *service-id* for local and remote service testing. The following information can be determined from **svc-ping**:

1. Local and remote service existence
2. Local and remote service state
3. Local and remote service type correlation
4. Local and remote customer association
5. Local and remote service-to-SDP bindings and state
6. Local and remote ingress and egress service label association

Unlike **sdp-ping**, only a single message will be sent per command; no count nor interval parameter is supported and round trip time is not calculated. A timeout value of 10 seconds is used before failing the request. The forwarding class is assumed to be Best-Effort Out-of-Profile

If no request is sent or a reply is not received, all remote information will be shown as N/A.

To terminate a **svc-ping** in progress, use the CLI break sequence <Ctrl-C>.

Upon request timeout, message response, request termination, or request error the following local and remote information will be displayed. Local and remote information will be dependent upon service existence and reception of reply.

Field	Description	Values
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 *service-id* 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 *service-id*. If the service does not exist locally, the administrative state will be Non-Existent.	Admin-Up
		Admin-Down
		Non-Existent
Local Service Oper State	The local operational state of *service-id*. If the service does not exist locally, the state will be N/A.	Oper-Up
		Oper-Down
		N/A

Field	Description	Values (Continued)
Remote Service Type	The remote type of service being tested. If *service-id* 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 *service-id*. 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 *service-id*. If the service does not exist, N/A is displayed.	*service-mtu* N/A
Remote Service MTU	The remote **service-mtu** for *service-id*. If the service does not exist remotely, N/A is displayed.	*remote-service-mtu* N/A
Local Customer ID	The local *customer-id* associated with *service-id*. If the service does not exist locally, N/A is displayed.	*customer-id* N/A
Remote Customer ID	The remote *customer-id* associated with *service-id*. If the service does not exist remotely, N/A is displayed.	*customer-id* 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.	*system-ip-address* 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 entered for the **svc-ping** command.	*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. **sdp-ping** should also fail.	*resp-ip-addr* N/A

Field	Description	Values (Continued)
Responders Expected Far-end Address	The expected source of the originator's *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* or the request is transmitted outside the *sdp-id*, N/A is displayed.	*resp-rec-tunnel-far-end-address* N/A
Originating SDP-ID	The *sdp-id* used to reach the **far-end** IP address if **sdp-path** is defined. The originating *sdp-id* must be bound to the *service-id* and terminate on the **far-end** IP address. If an appropriate originating *sdp-id* is not found, Non-Existent is displayed.	orig-sdp-id Non-Existent
Originating SDP-ID Path Used	Whether the Originating router used the originating *sdp-id* to send the **svc-ping** request. If a valid originating *sdp-id* is found, operational and has a valid egress service label, the originating router should use the *sdp-id* as the requesting path if **sdp-path** has been defined. If the originating router uses the originating *sdp-id* as the request path, Yes is displayed. If the originating router does not use the originating *sdp-id* as the request path, No is displayed. If the originating *sdp-id* is non-existent, N/A is displayed.	Yes No N/A
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 an originating *sdp-id* 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 *sdp-id*. If an originating *sdp-id* is not found, N/A is displayed.	Oper-Up Oper-Down N/A
Originating SDP-ID Binding Admin State	The local administrative state of the originating *sdp-id*s binding to *service-id*. If an *sdp-id* 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 *sdp-id*s binding to *service-id*. If an *sdp-id* is not bound to the service, N/A is displayed.	Oper-Up Oper-Down N/A
Responding SDP-ID	The *sdp-id* 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 *sdp-id* as the return path, but the appropriate responding *sdp-id* will be displayed. If a valid *sdp-id* return path is not found to the originating router that is bound to the *service-id*, Non-Existent is displayed.	*resp-sdp-id* Non-Existent

Field	Description	Values (Continued)
Responding SDP-ID Path Used	Whether the responding router used the responding *sdp-id* to respond to the **svc-ping** request. If the request was received via the originating *sdp-id* and a valid return *sdp-id* is found, operational and has a valid egress service label, the far-end router should use the *sdp-id* as the return *sdp-id*. If the far end uses the responding *sdp-id* as the return path, Yes is displayed. If the far end does not use the responding *sdp-id* as the return path, No is displayed. If the responding *sdp-id* is non-existent, N/A is displayed.	Yes No N/A
Responding SDP-ID Administrative State	The administrative state of the far-end *sdp-id* associated with the return path for *service-id*. When a return path is administratively down, Admin-Down is displayed. If the return *sdp-id* is administratively up, Admin-Up is displayed. If the responding *sdp-id* 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 *sdp-id* associated with the return path for *service-id*. 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 Binding Admin State	The local administrative state of the responder's *sdp-id* binding to *service-id*. If an *sdp-id* is not bound to the service, N/A is displayed.	Admin-Up Admin-Down N/A
Responding SDP-ID Binding Oper State	The local operational state of the responder's *sdp-id* binding to *service-id*. If an *sdp-id* 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 *sdp-id* to the far-end address that is bound to *service-id*. If the *sdp-id* signaling is off, *originator-vc-id* is 0. If the *originator-vc-id* does not exist, N/A is displayed.	*originator-vc-id* N/A
Responding VC-ID	The responder's VC-ID associated with the *sdp-id* to *originator-id* that is bound to *service-id*. If the *sdp-id* signaling is off or the service binding to *sdp-id* does not exist, *responder-vc-id* is 0. If a response is not received, N/A is displayed.	*responder-vc-id* N/A
Originating Egress Service Label	The originating service label (VC-Label) associated with the *service-id* for the originating *sdp-id*. If *service-id* does not exist locally, N/A is displayed. If *service-id* exists, but the egress service label has not been assigned, Non-Existent is displayed.	*egress-vc-label* N/A Non-Existent

Field	Description	Values (Continued)
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 *service-id* 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 *service-id* 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 *service-id* to the originating router. If *service-id* does not exist in the remote router, N/A is displayed. If *service-id* does exist remotely but the remote egress service label has not been assigned, Non-Existent is displayed.	*rec-vc-label* N/A Non-Existent
Responding Egress Service Label Source	The responder's egress service label source. If the responder's egress service label is manually defined, Manual is displayed. If the responder's egress service label is dynamically signaled, Signaled is displayed. If the *service-id* 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 *service-id* 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 *service-id* when sending to the originating router. If *service-id* does not exist locally, N/A is displayed. If *service-id* exists but an ingress service label has not been assigned, Non-Existent is displayed.	*ingress-vc-label* 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 *service-id* does not exist on the originator or the originators ingress service label has not been assigned, N/A is displayed.	Manual Signaled N/A

Field	Description	Values (Continued)
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 *service-id* 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 *service-id* when sending to the originating router. If *service-id* does not exist in the remote router, N/A is displayed. If *service-id* exists, but an ingress service label has not been assigned in the remote router, Non-Existent is displayed.	*resp-ingress-vc-label* N/A Non-Existent
Responders Ingress Label Source	The assigned ingress service label source on the remote router. If the ingress service label is manually defined on the remote router, Manual is displayed. If the ingress service label is dynamically signaled on the remote router, Signaled is displayed. If the *service-id* 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 *service-id* 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.

s**ervice** *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 7750 SR-Series to receive a reply message.

Values 1 — 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.

The following table indicates whether a message is sent and how the message is encapsulated based on the state of the service ID.

Local Service State	local-sdp Not Specified		local-sdp Specified	
	Message Sent	**Message Encapsulation**	**Message Sent**	**Message Encapsulation**
Invalid Local Service	Yes	Generic IP/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.

The following table indicates how the message response is encapsulated based on the state of the remote service ID.

Remote Service State	Message Encapsulation	
	remote-sdp Not Specified	remote-sdp Specified
Invalid Ingress Service Label	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
Invalid Service-ID	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
No Valid SDP-ID Bound on Service-ID	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
SDP-ID Valid But Down	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
SDP-ID Valid and Up, but No Service Label	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
SDP-ID Valid and Up, Egress Service Label, but VC-ID Mismatch	Generic IP/GRE OAM (PLP)	Generic IP/GRE OAM (PLP)
SDP-ID Valid and Up, Egress Service Label, but VC-ID Match	Generic IP/GRE OAM (PLP)	SDP Encapsulation with Egress Service Label (SLP)

Sample Output

```
*A: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** *service-id* [**sap** *sap-id*]
host-connectivity-verify subscriber *sub-ident-string* [**sla-profile** *sla-profile-name*]

Context oam

Description This command enables host connectivity verification checks.

Parameters **service** *service-id* **—** Specifies the service ID to diagnose or manage.

Values 1 — 2147483647

sap ***sap-id*** **—** Specifies the physical port identifier portion of the SAP definition. See <Link>Common CLI Command Descriptions on page 355 for command syntax.

sub-profile *sub-profile-name* **—** Specifies an existing subscriber profile name. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile *sla-profile-name* **—** Specifies an existing SLA profile name. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

vprn-ping

Syntax **vprn-ping** *service-id* **source** *ip-address* **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*]

Context <GLOBAL>
config>saa>test>type

Description This command performs a VPRN ping.

Parameters **service** *service-id* **—** The VPRN service ID to diagnose or manage.

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

source *ip-address* **—** The IP prefix for the source IP address in dotted decimal notation.

Values ipv4-address: 0.0.0.0 — 255.255.255.255
ipv6-address: x:x:x:x:x:x:x:x
x:x:x:x:x:x:d.d.d.d
x: [0..FFFF]H
d: [0..255]D

destination *ip-address* — The IP prefix for the destination IP address in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

size *octets* — The OAM request packet size in octets, expressed as a decimal integer.

Values 1 — 9198

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control — Specifies the response to come on the control plane.

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 100

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

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

profile {in | out} — The profile state of the MPLS echo request encapsulation.

Default out

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 <GLOBAL>
config>saa>test>type

Description Performs VPRN trace.

Parameters **service** *service-id* **—** The VPRN service ID to diagnose or manage.

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

source *src-ip* **—** The IP prefix for the source IP address in dotted decimal notation.

Values ipv4-address: 0.0.0.0 — 255.255.255.255
ipv6-address: x:x:x:x:x:x:x:x
x:x:x:x:x:x:d.d.d.d
x: [0..FFFF]H
d: [0..255]D

destination *dst-ip* **—** The IP prefix for the destination IP address in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255

size *octets* **—** The OAM request packet size in octets, expressed as a decimal integer.

min-ttl *vc-label-ttl* **—** The minimum TTL value in the VC label for the trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *vc-label-ttl* **—** The maximum TTL value in the VC label for the trace test, expressed as a decimal integer.

Default 4

Values 1 — 255

return-control **—** Specifies the OAM reply to a data plane OAM request be sent using the control plane instead of the data plane.

Default OAM reply sent using the data plane.

probe-count *send-count* **—** The number of OAM requests sent for a particular TTL value, expressed as a decimal integer.

Default 1

Values 1 — 10

interval *seconds* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

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 — 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

Sample Output

```
A:PE_1# oam vprn-trace 25 source 10.4.128.1 destination 10.16.128.0
TTL Seq Reply Node-id         Rcvd-on         Reply-Path RTT
----------------------------------------------------------------------------
[Send request TTL: 1, Seq. 1.]
1   1   1     10.128.0.4      cpm             In-Band    0ms
  Requestor 10.128.0.1 Route: 0.0.0.0/0
    Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
    Next Hops: [1] ldp tunnel
    Route Targets: [1]: target:65100:1
  Responder 10.128.0.4 Route: 10.16.128.0/24
    Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
    Next Hops: [1] ldp tunnel
    Route Targets: [1]: target:65001:100

[Send request TTL: 2, Seq. 1.]
2   1   1     10.128.0.3      cpm             In-Band    0ms
```

```
  Requestor 10.128.0.1 Route: 0.0.0.0/0
    Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
    Next Hops: [1] ldp tunnel
    Route Targets: [1]: target:65100:1
  Responder 10.128.0.3 Route: 10.16.128.0/24
    Vpn Label: 0 Metrics 0 Pref 0 Owner local
    Next Hops: [1] ifIdx 2 nextHopIp 10.16.128.0

[Send request TTL: 3, Seq. 1.]
[Send request TTL: 4, Seq. 1.]
...
----------------------------------------------------------------------------
A:PE_1#
```

VPLS MAC Diagnostics

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*] [**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 *service-id* —** The service ID of the service to diagnose or manage.

Values 1 — 2147483647

Values *service-id*: 1 — 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* — Specify an unused IP address in the same network that is associated with the VPLS.

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source-mac *ieee-address* — Specify the source MAC address that will be sent to the CPE. If not specified or set to 0, the MAC address configured for the CPM is used.

fc-name — The forwarding class of the MPLS echo request encapsulation.

Default be

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

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 — 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 — 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.

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 can be sent via the data plane or the control plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.
An **age** can be provided to age a particular OAM MAC using a specific interval. By default, OAM MAC addresses are not aged and can be removed with a **mac-purge** or with an FDB clear operation.

When split horizon group (SHG) is configured, the flooding domain depends on which SHG the packet originates from. The **target-sap** *sap-id* value dictates the originating SHG information.

Parameters **service** *service-id* **—** The Service ID of the service to diagnose or manage.

Values 1 — 2147483647

destination *ieee-address* **—** The MAC address to be populated.

flood — Sends the OAM MAC populate to all upstream nodes.

Default MAC populate only the local FIB.

age *seconds* **—** The age for the OAM MAC, expressed as a decimal integer.

Default The OAM MAC does not age.

Values 1 — 65535

force — Converts the MAC to an OAM MAC even if it currently another type of MAC.

Default Do not overwrite type.

t**arget-sap** *sap-id* **—** 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 place, 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. 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** *service-id* **target** *ieee-address* [**flood**] [**send-control**] [**register**]

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. When sending the MAC purge using the control plane, the packet is sent directly to the system IP address of the next hop.

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 can be sent via the data plane or the control plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

The **register** option reserves the MAC for OAM testing where it is no longer an active MAC in the FIB for forwarding, but it is retained in the FIB as a registered OAM MAC. Registering an OAM MAC prevents relearns for the MAC based on customer packets. Relearning a registered MAC can only be done through a **mac-populate** request. The originating SHG is always 0 (zero).

Parameters **service** *service-id* **—** The service ID of the service to diagnose or manage.

Values 1 — 2147483647

target *ieee-address* **—** The MAC address to be purged.

flood — Sends the OAM MAC purge to all upstream nodes.

Default MAC purge only the local FIB.

send-control — Send the mac-purge request using the control plane.

Default Request is sent using the data plane.

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 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 1 — 2147483647

destination *ieee-address* **—** The destination MAC address for the OAM MAC request.

size *octets* — The MAC OAM request packet size in octets, expressed as a decimal integer. The request payload is padded to the specified size with a 6 byte PAD header and a byte payload of 0xAA as necessary. If the octet size specified is less than the minimum packet, the minimum sized packet necessary to send the request is used.

Default No OAM packet padding.

Values 1 — 65535

ttl *vc-label-ttl* **—** The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM request sent using the data plane.

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the 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 — 10

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

mac-trace

Syntax **mac-trace service** *service-id* **destination** *ieee-address* [**size** *octets*] [**min-ttl** *vc-label-ttl*] [**max-ttl** *vc-label-ttl*] [**send-control**] [**return-control**] [**source** *ieee-address*] [**z-count** *probes-per-hop*] [**interval** *interval*] [**timeout** *timeout*]

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 Alcatel-Lucent OAM packets with increasing TTL values to determine the hop-by-hop route to a destination MAC.

In a MAC traceroute, the originating device creates a MAC ping echo request packet for the MAC to be tested with increasing values of the TTL. The echo request packet is sent through the control plane or data plane and awaits a TTL exceeded response or the echo reply packet from the device with the destination MAC. The devices that reply to the echo request packets with the TTL exceeded and the echo reply are displayed.

When a **source** *ieee-address* value is specified and the source MAC address is locally registered within a split horizon group (SHG), then this SHG membership will be used as if the packet originated from this SHG. In all other cases, SHG 0 (zero) will be used. Note that if the **mac-ping** is originated from a non-zero SHG, such packets will not go out to the same SHG.

If EMG is enabled, mac-trace will return only the first SAP in each chain.

Parameters **service** *service-id* **—** The Service ID of the service to diagnose or manage.

Values 1 — 2147483647

destination *ieee-address* **—** The destination MAC address to be traced.

size *octets* **—** The MAC OAM request packet size in octets, expressed as a decimal integer. The request payload is padded to the specified size with a 6 byte PAD header and a byte payload of 0xAA as necessary. If the octet size specified is less than the minimum packet, the minimum sized packet necessary to send the request is used.

Default No OAM packet padding.

Values 1 — 65535

min-ttl *vc-label-ttl* **—** The minimum TTL value in the VC label for the MAC trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *vc-label-ttl* **—** The maximum TTL value in the VC label for the MAC trace test, expressed as a decimal integer.

Default 4

Values 1 — 255

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 *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 — 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 — 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 — 10

IGMP Snooping Diagnostics

mfib-ping

Syntax **mfib-ping service** *service-id* **source** *src-ip* **destination** *mcast-address* [**size** *size*] [**ttl** *vc-label-ttl*] [**return-control**] [**interval** *interval*] [**send-count** *send-count*] [**timeout** *timeout*]

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.

Parameters **service** *service-id* — The service ID of the VPLS to diagnose or manage.

Values 1 — 2147483647

source *src-ip* — The source IP address for the OAM request.

destination *mcast-address* — 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 — 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 — 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 — 10

send-count *send-count* — The number of messages to send, expressed as a decimal integer. The count parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent.

The message interval value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *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 7750 SR will wait for a message reply after sending the next message request.

Upon the expiration of message timeout, the requesting 7750 SR 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 — 100

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.

Multicast FIB Connectivity Test Sample Output

```
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#
```

EFM Commands

efm

Syntax *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* — Specify the port ID in the slot/mda/port format.

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.

ETH-CFM OAM Commands

linktrace

Syntax **linktrace** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**ttl** *ttl-value*]

Context oam>eth-cfm

Default The command specifies to initiate a linktrace test.

Parameters *mac-address* — Specifies a unicast destination MAC address.

mep *mep-id* — Specifies the target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

ttl *ttl-value* — Specifies the TTL for a returned linktrace.

Values 0 — 255

loopback

Syntax **loopback** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**send-count** *send-count*] [**size** *data-size*] [**priority** *priority*]

Context oam>eth-cfm

Default The command specifies to initiate a loopback test.

Parameters *mac-address* — Specifies a unicast MAC address.

mep *mep-id* — Specifies target MAC address.

Values 1 — 8191

domain *md-index* — Specifies the MD index.

Values 1 — 4294967295

association *ma-index* — Specifies the MA index.

Values 1 — 4294967295

send-count *send-count* — Specifies the number of messages to send, expressed as a decimal integer. Loopback messages are sent back to back, with no delay between the transmissions.

Default 1

Values 1 — 5

size *data-size* **—** This is the size of the data portion of the data TLV. If 0 is specified no data TLV is added to the packet.

Values 0 — 1500

priority *priority* **—** Specifies a 3-bit value to be used in the VLAN tag, if present, in the transmitted frame.

Values 0 — 7

eth-test

Syntax *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*] [**data-length** *data-length*]

Context oam>eth-cfm

Description This command issues an ETH-CFM test.

Parameters *mac-address* — Specifies a unicast MAC address.

mep *mep-id* **—** Specifies target MAC address.

Values 1 — 8191

domain *md-index* **—** Specifies the MD index.

Values 1 — 4294967295

association *ma-index* **—** Specifies the MA index.

Values 1 — 4294967295

data-length *data-length* **—** Indicates the UDP data length of the echo reply, the length starting after the IP header of the echo reply.

Values 64 — 1500

Default 64

one-way-delay-test

Syntax **one-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]

Context oam>eth-cfm

Description This command issues an ETH-CFM one-way delay test.

Parameters *mac-address* — Specifies a unicast MAC address.

mep *mep-id* **—** Specifies target MAC address.

Values 1 — 8191

domain *md-index* **—** Specifies the MD index.

Values 1 — 4294967295

association *ma-index* **—** Specifies the MA index.

Values 1 — 4294967295

priority *priority* **—** Specifies the priority.

Values 0 — 7

Default The CCM and LTM priority of the MEP.

two-way-delay-test

Syntax **two-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]

Context oam>eth-cfm

Description This command issues an ETH-CFM two-way delay test.

Parameters *mac-address* — Specifies a unicast MAC address.

mep *mep-id* **—** Specifies target MAC address.

Values 1 — 8191

domain *md-index* **—** Specifies the MD index.

Values 1 — 4294967295

association *ma-index* **—** Specifies the MA index.

Values 1 — 4294967295

priority *priority* **—** Specifies the priority.

Values 0 — 7

Default The CCM and LTM priority of the MEP.

two-way-slm-test

Syntax **two-way-slm-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*] [**send-count** *send-count*] [**size** *data-size*] [**timeout** *timeout*] [**interval** *interval*]

Context oam>eth-cfm

Description This command configures an Ethernet CFM two-way SLM test in SAA.

mac-address — Specifies a unicast destination MAC address.

mep *mep-id* **—** Specifies the target MAC address.

Values 1 — 8191

domain *md-index* **—** Specifies the MD index.

Values 1 — 4294967295

association *ma-index* **—** Specifies the MA index.

Values 1 — 4294967295

priority *priority* **—** Specifies the priority.

Values 0—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. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

size *data-size* **—** This is the size of the data portion of the data TLV. If 0 is specified no data TLV is added to the packet.

Default 0

Values 0 — 1500

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

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** *name* [**owner** *test-owner*]
no test *name*

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 *name* — Identify the saa test name to be created or edited.

owner *test-owner* **—** Specifies the owner of an SAA operation up to 32 characters in length.

Values If a *test-owner* value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
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 *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

description

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

Context config>saa>test

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

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

continuous

Syntax [**no**] **continuous**

Context config>saa>test

Description 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 **saa** command.

The **no** form of the command disables the continuous running of the test. Use the **shutdown** command to disable the test.

jitter-event

Syntax **jitter-event rising-threshold** *threshold* [**falling-threshold** *threshold*] [**direction**]
no jitter-event

Context config>saa>test

Description Specifies that at the termination of an SAA test probe, the calculated jitter value is evaluated against the configured rising and falling jitter thresholds. SAA threshold events are generated as required.

Once the threshold (rising/falling) is crossed, it is disabled from generating additional events until the opposite threshold is crossed. If a falling-threshold is not supplied, the rising threshold will be re-enabled when it falls below the threshold after the initial crossing that generate the event.

The configuration of jitter event thresholds is optional.

Parameters **rising-threshold** *threshold* **—** Specifies a rising threshold jitter value. When the test run is completed, the calculated jitter value is compared to the configured jitter rising threshold. If the test run jitter value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483 milliseconds

falling-threshold *threshold* **—** Specifies a falling threshold jitter value. When the test run is completed, the calculated jitter value is compared to the configured jitter falling threshold. If the test run jitter value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483 milliseconds

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.
outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.
roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

latency-event

Syntax **latency-event rising-threshold** *threshold* [**falling-threshold** *threshold*] [**direction**]
no latency-event

Context config>saa>test

Description Specifies that at the termination of an SAA test probe, the calculated latency event value is evaluated against the configured rising and falling latency event thresholds. SAA threshold events are generated as required.

Once the threshold (rising/falling) is crossed, it is disabled from generating additional events until the opposite threshold is crossed. If a falling-threshold is not supplied, the rising threshold will be re-enabled when it falls below the threshold after the initial crossing that generate the event.

The configuration of latency event thresholds is optional.

Parameters **rising-threshold** *threshold* **—** Specifies a rising threshold latency value. When the test run is completed, the calculated latency value is compared to the configured latency rising threshold. If the test run latency value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483 milliseconds

falling-threshold *threshold* **—** Specifies a falling threshold latency value. When the test run is completed, the calculated latency value is compared to the configured latency falling threshold. If the test run latency value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483 milliseconds

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.
outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.
roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

loss-event

Syntax **loss-event rising-threshold** *threshold* [**falling-threshold** *threshold*] [**direction**]
no loss-event

Context config>saa>test

Description Specifies that at the termination of an SAA testrun, the calculated loss event value is evaluated against the configured rising and falling loss event thresholds. SAA threshold events are generated as required.

The configuration of loss event thresholds is optional.

Parameters **rising-threshold** *threshold* **—** Specifies a rising threshold loss event value. When the test run is completed, the calculated loss event value is compared to the configured loss event rising threshold. If the test run loss event value is greater than the configured rising threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 packets

falling-threshold *threshold* **—** Specifies a falling threshold loss event value. When the test run is completed, the calculated loss event value is compared to the configured loss event falling threshold. If the test run loss event value is greater than the configured falling threshold value then an SAA threshold event is generated. The SAA threshold event is tmnxOamSaaThreshold, logger application OAM, event #2101.

Default 0

Values 0 — 2147483647 packets

direction — Specifies the direction for OAM ping responses received for an OAM ping test run.

Values **inbound** — Monitor the value of jitter calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run.

outbound — Monitor the value of jitter calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run.
roundtrip — Monitor the value of jitter calculated for the round trip, two-way, OAM ping requests and replies for an OAM ping test run.

Default roundtrip

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

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 **type**
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** *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*] [**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 *service-id* —** The service ID of the service to diagnose or manage.

Values *service-id*: 1 — 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* — Specify an unused IP address in the same network that is associated with the VPLS.

ttl *vc-label-ttl* — The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source-mac *ieee-address* — Specify the source MAC address that will be sent to the CPE. If not specified or set to 0, the MAC address configured for the CPMCFM is used.

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

interval *interval* — The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

send-count ***send-count*** **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 255

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM request sent using the data plane.

dns

Syntax **dns target-addr** *dns-name* **name-server** *ip-address* [**source** *ip-address*] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context <GLOBAL>
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 — 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 — 120

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is

used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

eth-cfm-linktrace

Syntax **eth-cfm-linktrace** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**ttl** *ttl-value*] [**fc** {*fc-name*} [**profile** {**in**|**out**}]] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context config>saa>test>type

Description This command configures a CFM linktrace test in SAA.

Parameters *mac-address* — Specifies a unicast destination MAC address.

mep *mep-id* **—** Specifies the target MAC address.

Values 1 — 8191

domain *md-index* **—** Specifies the MD index.

Values 1 — 4294967295

association *ma-index* **—** Specifies the MA index.

Values 1 — 4294967295

ttl *ttl-value* **—** Specifies the maximum number of hops traversed in the linktrace.

Default 64

Values 1— 255

fc *fc-name* **—** The **fc** parameter is used to indicate the forwarding class of the 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. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 10

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

eth-cfm-loopback

Syntax **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*][**timeout** *timeout*] [**interval** *interval*]

Context config>saa>test>type

Description This command configures an Ethernet CFM loopback test in SAA.

mac-address — Specifies a unicast destination MAC address.

mep *mep-id* **—** Specifies the target MAC address.

Values 1 — 8191

domain *md-index* **—** Specifies the MD index.

Values 1 — 4294967295

association *ma-index* **—** Specifies the MA index.

Values 1 — 4294967295

size *data-size* **—** This is the size of the data portion of the data TLV. If 0 is specified no data TLV is added to the packet.

Default 0

Values 0 — 1500

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.

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. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

eth-cfm-two-way-delay

Syntax **eth-cfm-two-way-delay** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**fc** {*fc-name*} [**profile** {**in**|**out**}]] [**send-count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context config>saa>test>type

Description This command configures an Ethernet CFM two-way delay test in SAA.

mac-address — Specifies a unicast destination MAC address.

mep *mep-id* **—** Specifies the target MAC address.

Values 1 — 8191

domain *md-index* **—** Specifies the MD index.

Values 1 — 4294967295

association *ma-index* **—** Specifies the MA index.

Values 1 — 4294967295

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.

Default nc

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

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

eth-cfm-two-way-slm

Syntax **eth-cfm-two-way-delay** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**fc** {*fc-name*}] [**send-count** *send-count*] [**size** *data-size*] [**timeout** *timeout*] [**interval** *interval*]

Context config>saa>test>type

Description This command configures an Ethernet CFM two-way SLM test in SAA.

mac-address — Specifies a unicast destination MAC address.

mep *mep-id* **—** Specifies the target MAC address.

Values 1 — 8191

domain *md-index* **—** Specifies the MD index.

Values 1 — 4294967295

association *ma-index* **—** Specifies the MA index.

Values 1 — 4294967295

fc *fc-name* **—** The **fc** parameter is used to indicate the forwarding class of the 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. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

size *data-size* **—** This is the size of the data portion of the data TLV. If 0 is specified no data TLV is added to the packet.

Default 0

Values 0 — 1500

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is

used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request. The **timeout** value must be less than the **interval**.

Default 5

Values 1 — 10

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*] [**timeout** *timeout*]

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:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

dns-name — The DNS name of the far-end device to which to send the **svc-ping** request message, expressed as a character string up to 63 characters maximum.

Values ipv6-address: x:x:x:x:x:x:x:x[-interface]
x:x:x:x:x:x:d.d.d.d[-interface]
x: [0 — FFFF]H
d: [0 — 255]D
interface (32 chars max, mandatory for link local addresses)

rapid — Packets will be generated as fast as possible instead of the default 1 per second.

detail — Displays detailed information.

ttl *time-to-live* **—** The TTL value for the MPLS label, expressed as a decimal integer.

Values 1 — 128

tos *type-of-service* **—** Specifies the service type.

Values 0 — 255

size *bytes* **—** The request packet size in bytes, expressed as a decimal integer.

Values 0 — 16384

pattern *pattern* **—** The date portion in a ping packet will be filled with the pattern value specified. If not specified, position info will be filled instead.

Values 0 — 65535

source *ip-address|dns-name* **—** Specifies the IP address to be used.

Values
ipv4-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x:x
x:x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D
dns-name: 128 characters max

interval *seconds* **—** This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

next-hop *ip-address* **—** Only displays static routes with the specified next hop IP address.

Values
ipv4-address: a.b.c.d (host bits must be 0)
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 — FFFF]H
d: [0 — 255]D

interface ***interface-name*** **—** The name used to refer to the interface. The name must already exist in the **config>router>interface** context.

bypass-routing — Specifies whether to send the ping request to a host on a directly attached network bypassing the routing table.

count *requests* **—** Specifies the number of times to perform an OAM ping probe operation. Each OAM echo message request must either timeout or receive a reply before the next message request is sent.

Values 1 — 100000

Default 5

do-not-fragment — Sets the DF (Do Not Fragment) bit in the ICMP ping packet.

router *router-instance* **—** Specifies the router name or service ID.

Values
router-name: Base , management
service-id: 1 — 2147483647

Default Base

service-name *service-name* **—** Specifies the service name as an integer or string.

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

timeout *timeout* **—** Overrides the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

icmp-trace

Syntax **icmp-trace** [*ip-address* **|** *dns-name*] [**ttl** *time-to-live*] [**wait** *milli-seconds*] [**tos** *type-of-service*] [**source** *ip-address*] [**tos** *type-of-service*] [**router** *router-instance* **| service-name** *service-name*]

Context config>saa>test>type

Description 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:x:d.d.d.d	
	x:	[0 — FFFF]H
	d:	[0 — 255]D

dns-name — The DNS name of the far-end device to which to send the **svc-ping** request message, expressed as a character string to 63 characters maximum.

ttl *time-to-live* **—** The TTL value for the MPLS label, expressed as a decimal integer.

Values 1 — 255

wait *milliseconds* **—** The time in milliseconds to wait for a response to a probe, expressed as a decimal integer.

Default 5000

Values 1 — 60000

tos *type-of-service* **—** Specifies the service type.

Values 0 — 255

source *ip-address* **—** Specifies the IP address to be used.

Values

ipv4-address:	a.b.c.d	
ipv6-address:	x:x:x:x:x:x:x:x	
	x:x:x:x:x:x:d.d.d.d	
	x:	[0 — FFFF]H
	d:	[0 — 255]D

router *router-instance* **—** Specifies the router name or service ID.

Values *router-name*: Base , management
service-id: 1 — 2147483647

Default Base

lsp-ping

Syntax **lsp-ping** {{[*lsp-name*] [**path** *path-name*]} **|** {**prefix** *ip-prefix/mask*}} [**src-ip-address** *ip-addr*] [**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**]

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.

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 ping request.

Default The active LSP path.

Values Any path name associated with the LSP.

prefix *ip-prefix/mask* **—** Specifies the address prefix and subnet mask of the destination node.

src-ip-address *ip-addr* — Specifies the source IP address. This option is used when an OAM packet must be generated from a different address than the node's system interface address. An example is when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next-hop is set to an address other than the system interface address.

Values ipv4-address: a.b.c.d

fc *fc-name* — The **fc** parameter is used to indicate the forwarding class of the MPLS echo request 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 7750 SR 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.

Default be

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

src-ip-address *ip-addr* **—** This parameter specifies the source IP address. This parameter is used when an OAM packet must be generated from a different address than the node's system interface address. For example, when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next-hop is set to an address other than the system interface address.

Values ipv4-address: a.b.c.d

profile {**in** | **out**} **—** The profile state of the MPLS echo request encapsulation.

Default out

size *octets* **—** The MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 68 — The system sends the minimum packet size, depending on the type of LSP. No padding is added.

Values 84 — 65535

ttl *label-ttl* **—** The TTL value for the MPLS label, expressed as a decimal integer.

Default 255

Values 1 — 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.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

path-destination *ip-address* **—** Specifies the IP address of the path destination.

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:x (eight 16-bit pieces)		
	x:x:x:x:x:x:d.d.d.d		
	x:	[0 — FFFF]H	
	d:	[0 — 255]D	

lsp-trace

Syntax **lsp-trace** {{[*lsp-name]* [**path** *path-name*]} **|** {**prefix** *ip-prefix/mask*}} [**src-ip-address** *ip-addr*] [**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**]

Context oam
config>saa>test>type

Description This command displays the hop-by-hop path for an LSP.

The **lsp-trace** command performs an LSP traceroute using the protocol and data structures defined in the IETF draft (draft-ietf-mpls-lsp-ping-02.txt).

The LSP 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.

In an LSP traceroute, 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 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 pathname along which to send the LSP trace request.

Default The active LSP path.

Values Any path name associated with the LSP.

prefix *ip-prefix/mask* **—** Specifies the address prefix and subnet mask of the destination node.

size *octets* — The MPLS echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 68 — The system sends the minimum packet size, depending on the type of LSP. No padding is added.

Values 84 — 65535

src-ip-address *ip-addr* — Specifies the source IP address. This option is used when an OAM packet must be generated from a different address than the node's system interface address. An example is when the OAM packet is sent over an LDP LSP and the LDP LSR-ID of the corresponding LDP session to the next-hop is set to an address other than the system interface address.

Values ipv4-address: a.b.c.d

min-ttl *min-label-ttl* **—** The minimum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 1

Values 1 — 255

max-ttl *max-label-ttl* **—** The maximum TTL value in the MPLS label for the LDP treetrace test, expressed as a decimal integer.

Default 30

Values 1 — 255

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 — 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.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the 7750 SR will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A 'request timeout' message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 3

Values 1 — 10

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

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 7750 SR 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 7750 SR.

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

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 can be sent via the control plane or the data plane. The **send-control** option specifies the request be sent using the control plane. If **send-control** is not specified, the request is sent using the data plane.

A **mac-ping** is forwarded along the flooding domain if no MAC address bindings exist. If MAC address bindings exist, then the packet is forwarded along those paths, provided they are active. A response is generated only when there is an egress SAP binding for that MAC address or if the MAC address is a “local” OAM MAC address associated with the device’s control plan.

A **mac-ping** reply can be sent using the data plane or the control plane. The **return-control** option specifies the reply be sent using the control plane. If **return-control** is not specified, the request is sent using the data plane.

A **mac-ping** with data plane reply can only be initiated on nodes that can have an egress MAC address binding. A node without a FIB and without any SAPs cannot have an egress MAC address binding, so it is not a node where replies in the data plane will be trapped and sent up to the control plane.

A control plane request is responded to via a control plane reply only.

By default, MAC OAM requests are sent with the system or chassis MAC address as the source MAC. The **source** option allows overriding of the default source MAC for the request with a specific MAC address.

When a **source** *ieee-address* value is specified and the source MAC address is locally registered within a split horizon group (SHG), then this SHG membership will be used as if the packet originated from this SHG. In all other cases, SHG 0 (zero) will be used. Note that if the **mac-trace** is originated from a non-zero SHG, such packets will not go out to the same SHG.

If EMG is enabled, mac-ping will return only the first SAP in each chain.

Parameters **service** *service-id* **—** The service ID of the service to diagnose or manage.

Values *service-id*: 1 — 2147483647
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.

Default No OAM packet padding.

Values 1 — 65535

ttl *vc-label-ttl* **—** The TTL value in the VC label for the OAM MAC request, expressed as a decimal integer.

Default 255

Values 1 — 255

send-control — Specifies the MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM request sent using the data plane.

return-control — Specifies the MAC OAM reply to a data plane MAC OAM request be sent using the control plane instead of the data plane.

Default MAC OAM reply sent using the data plane.

source *src-ieee-address* **—** The source MAC address from which the OAM MAC request originates. By default, the system MAC address for the chassis is used.

Default The system MAC address.

Values Any unicast MAC value.

fc *fc-name* **—** The **fc** parameter is used to test the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

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

interval *interval* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second where the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

timeout *timeout* **—** The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default **timeout** value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

sdp-ping

Syntax **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**timeout** *seconds*] [**interval** *seconds*] [**size** *octets*] [**send-count** *send-count*]

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 *service-id* is not specified.

For round trip connectivity testing, the **resp-sdp** keyword must be specified. If **resp-sdp** is not specified, a uni-directional SDP test is performed.

To terminate an **sdp-ping** in progress, use the CLI break sequence <Ctrl-C>.

An **sdp-ping** response message indicates the result of the **sdp-ping** message request. When multiple response messages apply to a single SDP echo request/reply sequence, the response message with the highest precedence will be displayed. The following table displays the response messages sorted by precedence.

Result of Request	Displayed Response Message	Precedence
Request timeout without reply	Request Timeout	1
Request not sent due to non-existent *orig-sdp-id*	Orig-SDP Non-Existent	2

Result of Request	Displayed Response Message	Precedence
Request not sent due to administratively down *orig-sdp-id*	Orig-SDP Admin-Down	3
Request not sent due to operationally down *orig-sdp-id*	Orig-SDP Oper-Down	4
Request terminated by user before reply or timeout	Request Terminated	5
Reply received, invalid *origination-id*	Far End: Originator-ID Invalid	6
Reply received, invalid *responder-id*	Far End: Responder-ID Error	7
Reply received, non-existent *resp-sdp-id*	Far End: Resp-SDP Non-Existent	8
Reply received, invalid *resp-sdp-id*	Far End: Resp-SDP Invalid	9
Reply received, *resp-sdp-id* down (admin or oper)	Far-end: Resp-SDP Down	10
Reply received, No Error	Success	11

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 — 17407

resp-sdp *resp-sdp-id* **—** Optional parameter is used to specify the return SDP-ID to be used by the far-end 7750 SR for the message reply for round trip SDP connectivity testing. If *resp-sdp-id* does not exist on the far-end 7750 SR, terminates on another 7750 SR different than the originating 7750 SR, 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.

Default null. Use the non-SDP return path for message reply.

Values 1 — 17407

fc *fc-name* **—** The **fc** parameter is used to indicate the forwarding class of the SDP encapsulation. The actual forwarding class encoding is controlled by the network egress DSCP or LSP-EXP mappings.

The DSCP or LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end 7750 SR that receives the message request. The egress mappings of the egress network interface on the far-end 7750 SR 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 7750 SR. This is displayed in the response message output upon receipt of the message reply.

Default be

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

profile {**in** | **out**} **—** The profile state of the SDP 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.

Default 5

Values 1 — 10

interval *seconds* **—** The **interval** parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the **interval** is set to 1 second, and the **timeout** value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

size *octets* **—** The **size** parameter in octets, expressed as a decimal integer. This parameter is used to override the default message size for the **sdp-ping** request. Changing the message size is a method of checking the ability of an SDP to support a **path-mtu**. The size of the message does not include the SDP encapsulation, VC-Label (if applied) or any DLC headers or trailers.

When the OAM message request is encapsulated in an IP/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.

Default 40

Values 40 — 9198

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The **count** parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message **interval** value must be expired before the next message request is sent.

Default 1

Values 1 — 100

SpecialCases Single Response Connectivity Tests — A single response sdp-ping test provides detailed test results.

Upon request timeout, message response, request termination, or request error the following local and remote information will be displayed. Local and remote information will be dependent upon SDP-ID existence and reception of reply.

Field	Description	Values
Request Result	The result of the **sdp-ping** request message.	Sent - Request Timeout Sent - Request Terminated Sent - Reply Received Not Sent - Non-Existent Local SDP-ID Not Sent - Local SDP-ID Down
Originating SDP-ID	The originating SDP-ID specified by **orig-sdp**.	*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 *orig-sdp-id* 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 *orig-sdp-id* does not exist, N/A will be displayed.	Oper-Up Oper-Down N/A
Originating SDP-ID Path MTU	The local **path-mtu** for *orig-sdp-id*. If *orig-sdp-id* 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 **sdp-ping** request. If **resp-sdp** is not specified, the responding router will not use an SDP-ID as the return path and N/A will be displayed.	*resp-sdp-id* N/A
Responding SDP-ID Path Used	Displays whether the responding 7750 SR used the responding *sdp-id* to respond to the **sdp-ping** request. If *resp-sdp-id* 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
Responding SDP-ID Administrative State	The administrative state of the responding *sdp-id*. When *resp-sdp-id* is administratively down, Admin-Down will be displayed. When *resp-sdp-id* is administratively up, Admin-Up will be displayed. When *resp-sdp-id* exists on the far-end 7750 SR but is not valid for the originating router, Invalid is displayed. When *resp-sdp-id* does not exist on the far-end router, Non-Existent is displayed. When **resp-sdp** is not specified, N/A is displayed.	Admin-Down Admin-Up Invalid Non-Existent N/A

Field	Description	Values
Responding SDP-ID Operational State	The operational state of the far-end *sdp-id* associated with the return path for *service-id*. 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 *resp-sdp-id*. If *resp-sdp-id* 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 *orig-sdp-id*.	*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 7750 SR-Series 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
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

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* [**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*] [**timeout** *timeout*] [**interval** *interval*] [**ttl** *vc-label-ttl*]

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 *sdpid:vcid* parameter. However, if the ping is across two or more segments, at least the *sdpId:vcId*, **src-ip-address** *ip-addr*, **dst-ip-address** *ip-addr*, **ttl** *vc-label-ttl*and **pw-id** *pw-id* parameters are used where:

- The *src-ip-address* is system IP address of the router preceding the destination router.
- The *pwid* is actually the VC ID of the last pseudowire segment.
- The *vc-label-ttl* must have a value equal or higher than the number of pseudowire segments.

Note that VCCV ping is a multi-segment pseudowire. For a single-hop pseudowire, only the peer VCCV CC bit of the control word is advertised when the control word is enabled on the pseudowire. VCCV ping on multi-segment pseudowires require that the control word be enabled in all segments of the VLL.

If the control word is not enabled on spoke SDP it will not be signaled peer VCCV CC bits to the far end, consequently VCCV ping cannot be successfully initiated on that specific spoke SDP.

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* — The VC ID of the pseudowire being tested must be indicated with this parameter. The VC ID needs to exist on the local router and the far-end peer needs to indicate that it supports VCCV to allow the user to send vccv-ping message.

Values 1 — 17407:1 — 4294967295

src-ip-address *ip-addr* **—** Specifies the source IP address.

Values ipv4-address: a.b.c.d

dst-ip-address *ip-address* **—** Specifies the destination IP address.

Values ipv4-address: a.b.c.d

pw-id *pw-id* **—** Specifies the pseudowire ID to be used for performing a **vccv-ping** operation. The pseudowire ID is a non-zero 32-bit connection ID required by the FEC 128, as defined in RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*.

reply-mode {**ip-routed** | **control-channel**} **—** The reply-mode parameter indicates to the far-end how to send the reply message.The option control-channel indicates a reply mode in-band using vccv control channel.

Default control-channel

fc *fc-name* **—** The **fc** parameter is used to indicate the forwarding class of the MPLS echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal 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.

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

timeout *seconds* **—** The timeout parameter, in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A ‘request timeout’ message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 5

Values 1 — 10

interval *seconds* **—** The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the interval is set to 1 second, and the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 10

size *octets* **—** The VCCV ping echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 88

Values 88 — 9198

send-count *send-count* **—** The number of messages to send, expressed as a decimal integer. The count parameter is used to override the default number of message requests sent. Each message request must either timeout or receive a reply before the next message request is sent. The message interval value must be expired before the next message request is sent.

Default 1

Values 1 — 100

ttl *vc-label-ttl* **—** Specifies the time-to-live value for the vc-label of the echo request message. The outer label TTL is still set to the default of 255 regardless of this value.

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** *sdp-id:vc-id* [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**reply-mode** *ip-routed*|*control-channel*] [**probe-count** *probes-per-hop*] [**timeout** *timeout*] [**interval** *interval*] [**min-ttl** *min-vc-label-ttl*] [**max-ttl** *max-vc-label-ttl*] [**max-fail** *no-response-count*] [**detail**]

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 Addressí 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.

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

sdpid:vcid — The VC ID of the pseudowire being tested must be indicated with this parameter. The VC ID needs to exist on the local 7750 SR and the far-end peer needs to indicate that it supports VCCV to allow the user to send vccv-ping message.

Values 1-17407:1 — 4294967295

reply-mode {***ip-routed*** | ***control-channel***} **—** The reply-mode parameter indicates to the far-end how to send the reply message. The option control-channel indicates a reply mode in-band using vccv control channel.

Note that when a VCCV trace message is originated from an S-PE node, the user should used the IPv4 reply mode as the replying node does not know how to set the TTL to reach the sending S-PE node. If the user attempts this, a warning is issued to use the ipv4 reply mode.

Default control-channel

fc *fc-name* [**profile** {**in** | **out**} **—** The fc and profile parameters are used to indicate the forwarding class of the VCCV trace echo request packets. The actual forwarding class encoding is controlled by the network egress LSP-EXP mappings.

The LSP-EXP mappings on the receive network interface controls the mapping back to the internal forwarding class used by the far-end router that receives the message request. The egress mappings of the egress network interface on the far-end router controls the forwarding class markings on the return reply message. The LSP-EXP mappings on the receive network interface controls the mapping of the message reply back at the originating router.

fc-name — The forwarding class of the VCCV trace echo request encapsulation.

Default be

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

profile {**in** | **out**} **—** The profile state of the VCCV trace echo request encapsulation.

Default out

size *octets* **—** The VCCV ping echo request packet size in octets, expressed as a decimal integer. The request payload is padded with zeroes to the specified size.

Default 88

Values 88 — 9198

probe-count *probes-per-ho***p —** The number of VCCV trace echo request messages to send per TTL value.

Default 1

Values 1 — 10

timeout *timeout* **—** The timeout parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. A request timeout message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Default 3

Values 1 — 60

interval *interval* **—** The interval parameter in seconds, expressed as a decimal integer. This parameter is used to override the default request message send interval and defines the minimum amount of time that must expire before the next message request is sent.

If the interval is set to 1 second, and the timeout value is set to 10 seconds, then the maximum time between message requests is 10 seconds and the minimum is 1 second. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Default 1

Values 1 — 255

min-ttl *min-vc-label-ttl* **—** The TTL value for the VC label of the echo request message for the first hop of the MS-PW for which the results are to be displayed. This is expressed as a decimal integer. Note that the outer label TTL is still set to the default of 255 regardless of the value of the VC label.

Default 1

Values 1 — 255

max-ttl *max-vc-label-ttl* **—** The TTL value for the VC label of the echo request message for the last hop of the MS-PW for which the results are to be displayed. This is expressed as a decimal integer. Note that the outer label TTL is still set to the default of 255 regardless of the value of the VC label.

Default 8

Values 1 — 255

max-fail *no-response-count* **—** The maximum number of consecutive VCCV trace echo requests, expressed as a decimal integer that do not receive a reply before the trace operation fails for a given TTL value.

Default 5

Values 1 — 255

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#
```

OAM SAA Commands

saa

Syntax **saa** *test-name* [**owner** *test-owner*] {**start** | **stop**} [**no-accounting**]

Context oam

Description Use this command to start or stop an SAA test.

test-name — Name of the SAA test. The test name must already be configured in the **config>saa>test** context.

owner *test-owner* **—** Specifies the owner of an SAA operation up to 32 characters in length.

Values If a *test-owner* value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

start **—** This keyword starts the test. A test cannot be started if the same test is still running.

A test cannot be started if it is in a shut-down state. An error message and log event will be generated to indicate a failed attempt to start an SAA test run. A test cannot be started if it is in a continous 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 continous 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.

LDP Treetrace Commands

ldp-treetrace

Syntax **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*]]

Context oam

Description This command enables the context to configure LDP treetrace parameters to perform Alcatel-Lucent OAM tree trace test operations manually.

Parameters **prefix** *ip-prefix/mask* **—** Specifies the address prefix and subnet mask of the destination node.

max-ttl *max-label-ttl* — The maximum TTL value in the MPLS label for the LSP trace test, expressed as a decimal integer.

Default 30

Values 1 — 255

max-paths *max-paths* — The maximum number of paths for a ldp-treetrace test, expressed as a decimal integer.

Default 128

Values 1 — 255

timeout *timeout* — The **timeout** parameter in seconds, expressed as a decimal integer. This value is used to override the default timeout value and is the amount of time that the router will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting router assumes that the message response will not be received. Any response received after the request times out will be silently discarded.

Default 3

Values 1 — 60

fc *fc-name* — The **fc** and **profile** parameters are 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 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.

Default be

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

profile *profile* — The profile state of the MPLS echo request encapsulation.

Default out

Values in, 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.

Default 5

Values 1 — 255

ldp-treetrace

Syntax [**no**] **ldp-treetrace**

Context config>test-oam

Description This command enables the context to configure LDP treetrace parameters to perform OAM tree trace test operations manually**.**

The **no** form of the command disables the LDP treetrace parameters.

fc

Syntax **fc** *fc-name* [**profile** {**in** | **out**}]
no fc

Context config>test-oam>ldp-treetrace

Description This command configures forwarding class name and profile parameters. These parameters 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 SR7450 ESS7710 SR7210 SAS M that receives the message request. The egress mappings of the egress network interface on the far-end 7750 SR7450 ESS7710 SR7210 SAS M 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.

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 enables the context to configure path discovery parameters.

interval

Syntax **interval** *minutes*
no interval

Context config>test-oam>ldp-treetrace>path-discovery

Description This command configures the time to wait before repeating the LDP tree auto discovery process.

Default 60

Parameters *minutes* — Specifies the number of minutes to wait before repeating the LDP tree auto discovery process.

Values 60 — 1440

max-path

Syntax **max-path** *max-paths*

Context config>test-oam>ldp-treetrace>path-discovery

Description This command configures specifies the maximum number of paths that can be discovered for a selected IP address FEC.

Default 128

Parameters *max-paths* — Specifies the tree discovery maximum path.

Values 1 — 128

max-ttl

Syntax **max-ttl** *ttl-value*

Context config>test-oam>ldp-treetrace>path-discovery

Description This command configures the maximum label time-to-live value for an LSP trace request during the tree discovery.

Default 30

Parameters *ttl-value* — Specifies the maximum label time-to-live value for an LSP trace request during the tree discovery.

Values 1 — 255

policy-statement

Syntax **policy-statement** *policy-name* [...(up to 5 max)]

Context config>test-oam>ldp-treetrace>path-discovery

Description This command specifies policies to filter LDP imported address FECs.

Default no policy-statement

Parameters *policy-name* — Specifies the route policy name to filter LDP imported address FECs. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. The specified policy name(s) must already be defined.

retry-count

Syntax **retry-count** *retry-count*

Context config>oam-test>ldp-treetrace>path-discovery
config>oam-test>ldp-treetrace>path-probing

Description This command configures the path probing maximum number of failures.

Default 3

Parameters *retry-count* — Specifies the maximum number of consecutive timeouts allowed before failing a path probe (ping).

Values 1 — 255

timeout

Syntax **timeout** *timeout*
no timeout

Context config>test-oam>ldp-treetrace>path-discovery

Description This command 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 30

Parameters *timeout* — Specifies the timeout parameter, in seconds, within a range of 1 to 60, expressed as a decimal integer.

path-probing

Syntax **path-probing**

Context config>test-oam>ldp-treetrace

Description This command enables the context to configure path probing paramters.

interval

Syntax **interval** *minutes*
no interval

Context config>test-oam>ldp-treetrace>path-probing

Description This command configures the number of minutes to wait before repeating probing (pinging) a discovered path.

Default 1

Parameters *minutes* — Specifies the number of minutes to probe all active ECMP paths for each LSP

Values 1 — 60

retry-count

Syntax **retry-count** *retry-count*

Context config>oam-test>ldp-treetrace>path-discovery
config>oam-test>ldp-treetrace>path-probing

Description This command configures the path probing maximum number of failures.

Default 3

Parameters *retry-count* — Specifies the maximum number of consecutive timeouts allowed before failing a path probe (ping).

Values 1 — 255

timeout

Syntax **timeout** *timeout*
no timeout

Context config>test-oam>ldp-treetrace>path-probing

Description This command 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 1

Parameters *timeout* — Specifies the timeout parameter, in minutes, with a range of 1 to 3 minutes, expressed as a decimal integer.

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 a 7x50 node 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 time stam’s *seconds* field holds the integral number of seconds since 1-Jan-1900 00:00:00 UTC. The time stamp's *microseconds* field contains a microseconds value in the range 0 — 999999. This setting is used to interoperate 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 stampss *seconds* field holds a Unix time, the integral number of seconds since 1-Jan-1970 00:00:00 UTC. The time stampss *microseconds* field contains a microseconds value in the range 0 — 999999. This setting is used to interoperate 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 SROS system running software earlier than R8.0 R4).

twamp

Syntax **twamp**

Context config>oam-test

Description This command enables TWAMP functionality.

Default TWAMP is disabled.

server

Syntax **retry-count** *retry-count*

Context config>test-oam>twamp

Description This command configures the node for TWAMP server functionality.

Default TWAMP is disabled.

prefix

Syntax **prefix** {*address/mask* **|** *address netmask*}
no prefix {*address/mask* **|** *address netmask*}

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 *address* — An IPv4 address in dotted decimal notation.

Values a.b.c.d

Default none

mask — The prefix length.

Values 0—32

Default none

retry-count — The netmask in dotted decimal notation.

Values a.b.c.d

Default none

max-conn-prefix

Syntax **max-conn-prefix** *count*
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 *count* — The maximum number of control connections.

Values 0—64

Default 32

max-conn-server

Syntax **max-conn-server** *count*
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 *count* — The maximum number of control connections.

Values 0—64

Default 32

inactivity-timeout

Syntax **inactivity-timeout** *seconds*
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 instructs the system to go with the default value of 1800 seconds.

Default no inactivity-timeout

Parameters *retry-count* — The duration of the inactivity timeout.

Values 0 — 3600

Default 1800

max-sess-prefix

Syntax **max-sess-prefix** *count*
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 means to go with a default value of 32.

Default no max-sess-prefix

Parameters *count* — The maximum number of concurrent test sessions.

Values 0— 128

Default 32

max-sess-server

Syntax **max-sess-server** *count*
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 *count* — The maximum number of concurrent test sessions.

Values 0— 128

Default 32

port

Syntax **port** *number*
no port

Context config>test-oam>twamp>server

Description This command configures the TCP port number used by the TWAMP server to listen for incoming connection requests from TWAMP clients.

The port number can be changed only when the server has been shutdown.

The no form of this command means to go with the default of 862.

Default no port

Parameters *number* — The TCP port number.

Values 1 — 65535

Default 862

Show Commands

saa

Syntax **saa** [*test-name*] [**owner** *test-owner*]

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 *test-name* — Enter the name of the SAA test for which the information needs to be displayed. The test name must already be configured in the **config>saa>test** context.

This is an optional parameter.

owner *test-owner* **—** Specifies the owner of an SAA operation up to 32 characters in length.

Values 32 characters maximum.

Default If a *test-owner* value is not specified, tests created by the CLI have a default owner “TiMOS CLI”.

Output **SAA Output —** The following table provides SAA field descriptions.

Label	Description
`Test Name`	Specifies the name of the test.
`Owner Name`	Specifies the owner of the test.
`Description`	Specifies the description for the test type.
`Accounting policy`	Specifies the associated accounting policy ID.
`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.

Label	Description (Continued)
`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.

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 asso-
ciation 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
...
===============================================================================
```

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.

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           10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.9           10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.15          10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.19          10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.24          10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.28          10.10.1.2           12.12.12.10         12/19/2006 05:11:01

……………..

127.1.0.252         10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  0                   EgressRtr
127.1.0.255         10.10.1.2           12.12.12.10         12/19/2006 05:11:01
              7     OK                  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 Rej   : 0
Tests Abort : 0
TstPktsRx   : 999                             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                              Oper State  : Up
Curr Conn   : 0                               Max Conn    : 32
Conn Reject : 0
Curr Sess   : 0                               Max Sess    : 32
Tests Done  : 0                               Tests Rej   : 0
Tests Abort : 0
TstPktsRx   : 0                               TstPktsTx   : 0
-------------------------------------------------------------------------------
Client          Sessions               Idle     TstPktsRx    TstPktsTx
                Curr/Done/Rej/Abort
-------------------------------------------------------------------------------
===============================================================================
```

eth-cfm

Syntax **eth-cfm**

Context show

Description This command enables the context to display CFM information.

association

Syntax **association** [*ma-index*] [**detail**]

Context show>eth-cfm

Description This command displays eth-cfm association information.

Parameters *ma-index* — Specifies the MA index.

Values 1— 4294967295

detail — Displays detailed information for the eth-cfm association.

Sample Output

```
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 <port-id> **[vlan** <qtag[.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 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 — 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.

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

===============================================================================
CFM SAP Stack Table
===============================================================================
Sap               Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
lag-1:100.100       3 Down          3          1  101 d0:0d:1e:00:01:01 ------
===============================================================================

===============================================================================
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
-------------------------------------------------------------------------------
1/1/10    0          0 Down         10          1    6 90:f4:01:01:00:0a --C---
===============================================================================

===============================================================================
CFM Facility LAG Stack Table
===============================================================================
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
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
CFM SDP Stack Table
===============================================================================
Sdp                Lvl Dir  Md-index   Ma-index   MepId  Mac-address     Defect
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================

===============================================================================
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.

Sample Output

```
*A:node-1# show eth-cfm domain

===============================================================================
CFM Domain Table
===============================================================================
Md-index   Level Name                                      Format
-------------------------------------------------------------------------------
1       4 test-1                     charString
2          5                                               none
25      7    AA:BB:CC:DD:EE:FF-1            macAddressAndUint
===============================================================================
```

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]
mep *mep-id* **domain** *md-index* **association** *ma-index* [**remote-mepid** *mep-id* **| all-remote-mepids**]
mep *mep-id* **domain** *md-index* **association** *ma-index* **eth-test-results** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **one-way-delay-test** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **two-way-delay-test** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **two-way-slm-test [remote-peer** *mac-address*]

Context show>eth-cfm

Description This command displays Maintenance Endpoint (MEP) information.

Parameters **domain** *md-index* **—** Displays the index of the MD to which the MP is associated, or 0, if none.

association ***ma-index*** **—** Displays the index to which the MP is associated, or 0, if none.

loopback — Displays loopback information for the specified MEP.

linktrace — Displays linktrace information for the specified MEP.

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.

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)
FngState           : fngReset                   ControlMep        : False
LowestDefectPri    : macRemErrXcon              HighestDefect     : none
Defect Flags       : None
Mac Address        : d0:0d:1e:00:01:01          ControlMep        : False
CcmLtmPriority     : 7
CcmTx              : 19886                      CcmSequenceErr    : 0
Fault Propagation  : disabled                   FacilityFault     : n/a
MA-CcmInterval     : 1                          MA-CcmHoldTime    : 0ms
```

```
Eth-1Dm Threshold  : 3(sec)                MD-Level          : 3
Eth-Ais:           : Enabled               Eth-Ais Rx Ais:   : No
Eth-Ais Tx Priorit*: 7                     Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                     Eth-Ais Tx Counte*: 388
Eth-Ais Tx Levels  : 5
Eth-Tst:           : Disabled

Redundancy:
    MC-LAG State   : active

CcmLastFailure Frame:
    None

XconCcmFailure Frame:
    None
===============================================================================
```

mip

Syntax **mip**

Context show>eth-cfm

Description This command displays SAPs/bindings provisioned for allowing the default MIP creation.

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
==========================================================================
```

system-config

Syntax **system-config**

Context show>eth-cfm

Description This command shows various system level configuration parameters. These global eth-cfm commands are those which are configured directly under the config>eth-cfm context.

Sample Output

```
# show eth-cfm system-config
===============================================================================
CFM System Configuration
===============================================================================
Redundancy
    MC-LAG Standby MEP Shutdown: true
    MC-LAG Hold-Timer          :   1 second(s)

Synthetic Loss Measurement
    Inactivity Timer           : 100 second(s)
===============================================================================
```

Clear Commands

saa

Syntax **saa-test** [*test-name* [**owner** *test-owner*]]

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 *test-name* — Name of the SAA test. The test name must already be configured in the **config>saa>test** context.

owner *test-owner* **—** Specifies the owner of an SAA operation up to 32 characters in length.

Default If a *test-owner* value is not specified, tests created by the CLI have a default owner "TiMOS CLI".

Debug Commands

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.

Tools Command Reference

Command Hierarchies

Configuration Commands

Tools Dump Commands

tools
- — **dump**
  - — **aps** *aps-id* [**clear**]
  - — **aps** **mc-aps-signaling** [**clear**]
  - — **aps** **mc-aps-ppp** [**clear**]
  - — **lag** **lag-id** *lag-id*
  - — **ldp-treetrace** {**prefix** *ip-prefix/mask*| **manual-prefix** *ip-prefix/mask*}[**path-destination** *ip-address*] [**trace-tree**]
  - — **persistence**
    - — **submgt** [**record** *record-key*]
    - — **summary**
  - — **ppp** *port-id*
  - — **redundancy**
    - — **multi-chassis**
      - — **mc-endpoint** **peer** *ip-address*
      - — **mc-ring**
      - — **mc-ring** **peer** *ip-address* [ring *sync-tag*]
      - — **srrp-sync-database** [**instance** *instance-id*] [**peer** *ip-address*]
      - — **sync-database** [**peer** *ip-address*] [**port** *port-id* | *lag-id*] [**sync-tag** *sync-tag*] [**application** *application*] [**detail**] [**type** *type*]
  - — **router** *router-instance*
    - — **dhcp**
      - — **group-if-mapping** [**clear**]
      - — **group-if-stats** [**clear**]
    - — **ldp**
      - — **fec** **prefix** *ip-prefix/mask*
      - — **fec** **vc-type** {**ethernet**|**vlan**} **vc-id** *vc-id*
      - — **instance**
      - — **interface** [*ip-int-name* | *ip-address*]
      - — **memory-usage**
      - — **peer** *ip-address*
      - — **session** [*ip-addr*[:*label-space*] [**connection**|**peer**|**adjacency**]
      - — **sockets**
      - — **timers**
    - — **mpls**
      - — **ftn** [**endpoint** *endpoint* | **sender** *sender* | **nexthop** *nexthop* | **lsp-id** *lsp-id* | **tunnel-id** *tunnel-id* | **label** *start-label end-label*]

— **ilm** [**endpoint** *endpoint* | **sender** *sender* | **nexthop** *nexthop* | **lsp-id** *lsp-id* | **tunnel-id** *tunnel-id* | **label** *start-label end-label*]
— **lspinfo** [*lsp-name*] [**detail**]
— **memory-usage**
— **te-lspinfo** [**endpoint** *ip-address*] [**sender** *ip-address*] [**lspid** *lsp-id*] [**detail**] [**p2p** | *p2p-tid tunnel-id*]
— **te-lspinfo** [**endpoint** *ip-address*] [**sender** *ip-address*] [**lspid** *lsp-id*] [**detail**] [**p2p** | *p2p-tid tunnel-id*]{ [**phops**] [**nhops**] [**s2l** *ip-address*] } }
— **ospf**
— **ospf3**
— **abr** [**detail**]
— **asbr** [**detail**]
— **bad-packet** *interface-name*
— **leaked-routes** [**summary** | **detail**]
— **memory-usage** [**detail**]
— **request-list** [**neighbor** *ip-address*] [**detail**]
— **request-list** **virtual-neighbor** *ip-address* **area-id** *area-id* [**detail**]
— **retransmission-list** [**neighbor** *ip-address*] [**detail**]
— **retransmission-list** **virtual-neighbor** *ip-address* **area-id** *area-id* [**detail**]
— **route-summary**
— **route-table** [**type**] [**detail**]
— **pim**
— **iom-failures** [**detail**]
— **rsvp**
— **psb** [**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** ***tun-nel-id***] [**lspid** ***lsp-id***]
— **rsb** [**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** ***tun-nel-id***] [**lspid** ***lsp-id***]
— **tcsb**[**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** ***tun-nel-id***] [**lspid** ***lsp-id***]
— **static-route** **ldp-sync-status**
— **web-rd**
— **http-client** [*ip-prefix/mask*]
— **service**
— **base-stats** [**clear**]
— **iom-stats** [**clear**]
— **l2pt-diags**
— **l2pt-diags** **clear**
— **l2pt-diags** **detail**
— **mc-endpoint** *mc-ep-id*
— **radius-discovery** [**svc-id** *service-id*]
— **vpls-fdb-stats** [**clear**]
— **vpls-mfib-stats** [**clear**]
— **system-resources** *slot-number*

Tools Perform Commands

tools
— **perform**
— **aps**
— **clear** *aps-id* {**protect** | **working**}
— **exercise** *aps-id* {**protect** | **working**}
— **force** *aps-id* {**protect** | **working**}
— **lockout** *aps-id*
— **request** *aps-id* {**protect** | **working**}
— **cron**
— **action**
— **stop** [*action-name*] [**owner** *action-owner*] [**all**]
— **tod**
— **re-evaluate**
— **customer** *customer-id* [**site** *customer-site-name*]
— **filter** *filter-type* [*filter-id*]
— **service** **id** *service-id* [**sap** *sap-id*]
— **tod-suite** *tod-suite-name*
— **lag**
— **clear-force** **all-mc**
— **clear-force** **lag-id** *lag-id* [**sub-group** *sub-group-id*]
— **clear-force** **peer-mc** *ip-address*
— **force** **all-mc** {**active** | **standby**}
— **force** **lag-id** *lag-id* [**sub-group** *sub-group-id*] {**active** | **standby**}
— **force** **peer-mc** *peer-ip-address* {**active** | **standby**}
— **log**
— **test-event**
— **router** [*router-instance*]
— **consistency**
— **isis**
— **ldp-sync-exit**
— **run-manual-spf**
— **mpls**
— **adjust-autobandwidth** [**lsp** *lsp-name* [**force** [**bandwidth** *mbps*]]]
— **cspf** **to** ***ip-addr*** [**from** *ip-addr*] [**bandwidth** *bandwidth*] [**include-bitmap** *bitmap*] [**exclude-bitmap** *bitmap*] [**hop-limit** *limit*] [**exclude-address** *excl-addr* [*excl-addr*...(up to 8 max)]] [**use-te-metric**] [**strict-srlg**] [**srlg-group** *grp-id*...(up to 8 max)] [**exclude-node** *excl-node-id* [*excl-node-id* ..(up to 8 max)]] [**skip-interface** *interface-name*] [**ds-class-type** *class-type*] [**cspf-reqtype** *req-type*] [**least-fill-min-thd** *thd*] [**setup-priority** *val*] [**hold-priority** *val*]
— **resignal** **lsp** *lsp-name* **path** *path-name* **delay** *minutes*
— **resignal** {**p2mp-lsp** *p2mp-lsp-name* **p2mp-instance** *p2mp-instance-name* | **p2mp-delay** *p2mp-minutes*}
— **trap-suppress** ***number-of-traps time-interval***
— **ospf** [*ospf-instance*]
— **ldp-sync-exit**
— **refresh-lsas** [*lsa-type*] [*area-id*]
— **run-manual-spf** *externals-only*
— **ospf3** [*ospf-instance*]
— **refresh-lsas** [*lsa-type*] [*area-id*]
— **run-manual-spf** *externals-only*

— **security**
— **authentication-server-check** *server-address ip-address* [**port** *port*] **user-name** *DHCP client user name* **password** *password* **secret** *key* [**source-address** *ip-address*] [**timeout** *seconds*] [**router** *router-instance*]
— **service**
— **egress-multicast-group** *group-name*
— **force-optimize**
— **eval-pw-template** *policy-id* [**allow-service-impact**]
— **id** *service-id*
— **endpoint** *endpoint-name*
— **force-switchover** *sdp-id:vc-id*
— **no force-switchover**
— **eval-pw-template** *policy-id* [**allow-service-impact**]
— **mcac sap** *sap-id* **recalc policy** *policy-name* [**bundle** *bundle-name*]
— **mcac sdp** *sdp-id:vc-id* **recalc policy** *policy-name* [**bundle** *bundle-name*]

tools
— **perform**
— **subscriber-mgmt**
— **edit-lease-state sap** *sap-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*]
— **edit-lease-state svc-id** *service-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*]
— **eval-lease-state** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ip** *ip-address*]
— **forcerenew svc-id** *service-id* {**ip** *ip-address*[/*mask*]>|**mac** *ieee-address*}
— **forcerenew** {**interface** *interface-name* | **sap** *sap-id*|**sdp** *sdp-id:vc-id*} [**ip** *ip-address*[/*mask*] |**mac** *ieee-address*]
— **re-ident-sub** *old-sub-ident-string* **to** *new-sub-ident-string*
— **remap-lease-state old-mac** *ieee-address* **mac** *ieee-address*
— **remap-lease-state sap** *sap-id* [**mac** *ieee-address*]

Tools Configuration Commands

Generic Commands

tools

Syntax **tools**

Context root

Description This command enables the context to enable useful tools for debugging purposes.

Default none

Parameters **dump —** Enables dump tools for the various protocols.

perform — Enables tools to perform specific tasks.

Dump Commands

dump

Syntax **dump** *router-name*

Context tools

Description The context to display information for debugging purposes.

Default none

Parameters *router-name* — Specify a router name, up to 32 characters in length.

Default Base

aps

Syntax **aps** *aps-id* [**clear**]
aps mc-aps-signaling [**clear**]
aps mc-aps-ppp [**clear**]

Context tools>dump>aps

Description This command displays Automated Protection Switching (APS) information.

Parameters **clear —** Removes all Automated Protection Switching (APS) operational commands.

mc-aps-signaling — Displays multi-chassis APS signaling information.

mc-aps-ppp — Displays multi-chassis APS PPP information.

Sample Output

```
*A:AS_SR7_2# tools dump aps aps-33

GrpId = 33, state = Running, mode:cfg/oper = Bi-directional/Bi-directional
   revert = 0, workPort: N/A, protPort: 2/1/1, activePort: working
   rxK1 = 0x0 (No-Req on Protect), physRxK1 = 0x0, rxK2 = 0x5
   txK1 = 0x0 (No-Req on Protect), physTxK1 = 0x0, txK2 = 0x5
   K1ReqToBeTxed = 0x0, K1ChanToBeTxed = 0x0, lastRxReq = 0xc
   MC-APS Nbr = 100.100.100.1 (Up), advIntvl = 10, hold = 30
   workPort: status = OK, Tx-Lais = None, sdCnt = 1, sfCnt = 1
     numSwitched = 1, switchSecs = 0, lastSwitched = 07/25/2007 08:00:12
     disCntTime = , alarms = , switchCmd = No Cmd
   protPort: status = OK, Tx-Lais = None, sdCnt = 1,  sfCnt = 0
     numSwitched = 1, switchSecs = 0, lastSwitched = 07/25/2007 08:03:39
     disCntTime = , alarms = ,  switchCmd = No Cmd
   GrpStatus: OK, mmCnt = 1, cmCnt = 1, psbfCnt = 1, feplfCnt = 2
   LocalSwitchCmd: priority = No-Req, portNum = 0
   RemoteSwitchCmd: priority = No-Req, portNum = 0
   Running Timers = mcAdvIntvl mcHold
   processFlag =  apsFailures = , sonet = Y
```

```
  DebugInfo: dmEv = 0, dmClrEv = 0, amEv = 1, amClrEv = 1
    cmEv = 1, cmClrEv = 1, psbfEv = 1, psbfClrEv = 1
    feplfEv = 2, feplfClrEv = 2, wtrEv = 0, psbfDetectEv = 0
    wSdEv = 1, wSfEv = 2, pSdEv = 1, pSfEv = 1
    portStatusEv = 8, rxK1Ev = 9, txLaisEv = 2, lastEvName = FeplClr
    CtlUpEv = 3, CtlDnEv = 2, wAct = 0, wDeAct = 0
Seq      Event  TxK1/K2 RxK1/K2  Dir    Active         Time
===     ======= ======= ======= =====   ======  ================
000     ProtAdd 0xc005  0x0000  Tx-->     Work  497 02:18:10.590
001     RxKByte 0xc005  0x6dea  Rx<--     Work  497 02:20:14.820
002     RxKByte 0xc005  0xc005  Rx<--     Work  497 02:21:30.970
003     RxKByte 0xc005  0x2005  Rx<--     Work  497 02:21:36.530
004      pSFClr 0x0005  0x2005  Tx-->     Work  497 02:21:40.590
005     RxKByte 0x0005  0x0005  Rx<--     Work  497 02:21:40.600
006     RxKByte 0x0005  0xc115  Rx<--     Work  497 02:25:22.840
007     RxKByte 0x2115  0xc115  Tx-->     Prot  497 02:25:22.840
008     RxKByte 0x2115  0xa115  Rx<--     Prot  000 00:00:47.070
009     RxKByte 0x2115  0x1115  Rx<--     Prot  000 00:00:47.560
010     RxKByte 0x2115  0xc005  Rx<--     Prot  000 00:00:57.010
011     RxKByte 0x2005  0xc005  Tx-->     Work  000 00:00:57.010
012     RxKByte 0x2005  0x0005  Rx<--     Work  000 00:01:06.170
013     RxKByte 0x0005  0x0005  Tx-->     Work  000 00:01:06.170
```

Sample Output

```
:AS_SR7_1# tools dump aps mc-aps-ppp

 pppmMcsModStarted = Yes
 pppmMcsDbgDoSync = Yes
 pppmMcsApsGrpHaAuditDone = Yes
 pppmMcsPostHaSyncedApsGrpId = 47
 pppmMcsMcApsChanCnt = 1280

 pppmMcsDbgRxPktCnt = 2560
 pppmMcsDbgRxPktNotProcessedCnt = 0
 pppmMcsDbgRxPktInvalidCnt = 0
 pppmMcsDbgInconsistentRxPktDropCnt = 0
 pppmMcsDbgInconsistentTxPktDropCnt = 1176
 pppmMcsDbgTxPktNotSentCnt = 0
 pppmMcsDbgTxPktSentCnt = 25
 pppmMcsDbgEvtDropCnt = 0
 pppmMcsDbgMemAllocErrCnt = 0
 pppmMcsDbgReTxCnt = 0
 pppmMcsDbgReTxExpCnt = 0
 pppmMcsDbgReReqCnt = 0

 pppmMcsStateAckQueueCnt (curr/peek) = 0/130
 pppmMcsStateReqQueueCnt (curr/peek) = 0/1280
 pppmMcsStateReReqQueueCnt (curr/peek) = 0/256
 pppmMcsStateTxQueueCnt (curr/peek) = 0/512
 pppmMcsStateReTxQueueCnt (curr/peek) = 0/130

 MC-APS Peer Info :
 --------------------

   Grp 13 Addr 100.100.100.2 - Up
   Grp 20 Addr 100.100.100.2 - Up
```

```
  Grp 35 Addr 100.100.100.2 - Up
  Grp 43 Addr 100.100.100.2 - Up
  Grp 47 Addr 100.100.100.2 - Up

Number of pppmMcs Evt Msgs dispatched:
  ctl_link_state : 0
  ctl_link_up_tmr : 0
  ctl_link_down_tmr : 0
  ha_audit_done : 0
```

Sample Output

```
*A:eth_aps_sr7# tools dump aps mc-aps-signaling

 MC-APS Control Debug Counters :
 -------------------------------
 Ctl Pkt Rx = 0
 Invalid Rx Ctl Pkt = 0
 Incompatible Rx Ctl Pkt = 0
 Nbr not Rx Ctl Pkt = 0
 Invalid Rx Ctl Pkt Tlv = 0
 Ctl Pkt Rx-ed before HaReady = 0
 Not sent Tx Ctl Pkt = 0

 MC-APS-LAG Debug Counters :
 ---------------------------
 Ctl Pkt Rx from IOM        = 0

 Not processed Rx Ctl Pkt   = 0
 Invalid Rx Ctl Pkt         = 0
 Incompatible Rx Ctl Pkt    = 0
 Rx Ctl Pkt queueing failed = 0

 Ctl Pkt Tx (direct)        = 0
 Ctl Pkt Tx (UDP socket)    = 0
 Not sent Tx Ctl Pkt        = 0

 Route Update               = 0
 Matched Route Update       = 0

 Msg Buf Alloc Failed       = 0

 MC-APS-LAG NbrRoute Entries :
 -------------------------------
 NbrAddr 1.1.1.1 NextHopAddr ::
   EgressIfIndex =  0
   EgressPortId =  Unknown
   app refCnt  =  1
   refCntTotal =  1
```

lag

Syntax **lag lag-id** *lag-id*

Context tools>dump

Description This tool displays LAG information.

Parameters *lag-id* — Specify an existing LAG id.

Values 1 — 200 (7750 SR-1: 1 — 64)

```
ALA-12>tools>dump# lag lag-id 1
Port state      : Ghost
Selected subgrp : 1
NumActivePorts  : 0
ThresholdRising : 0
ThresholdFalling: 0
IOM bitmask     : 0
Config MTU      : 1514
Oper. MTU       : 1514
Bandwidth       : 100000
ALA-12>tools>dump#
```

ldp-treetrace

Syntax **ldp-treetrace** {**prefix** *ip-prefix/mask*| **manual-prefix** *ip-prefix/mask*}[**path-destination** *ip-address*] [**trace-tree**]

Context tools>dump

Description This command displays TreeTrace information.

Parameters **prefix** *ip-prefix/mask* — Specifies the IP prefix and host bits.

Values		
	host bits:	must be 0
	mask:	0 — 32

Sample Output

Automated ldp-treetrace:

Note that the **tools dump ldp-treetrace prefix** command displays entries only if **ldp-treetrace** is enabled (**configure test-oam ldp-treetrace no shutdown**).

```
*A:Dut-B# tools dump ldp-treetrace prefix 10.20.1.6/32
    Discovered Paths:
    ===================
    Id   PathDst          EgrNextHop       ReplyRtrAddr     DiscoveryTime
           DiscoveryTtl     ProbeState       ProbeTmOutCnt    RtnCode
    ===  ================  ================  ================  ===================
    001       127.1.0.255       10.10.41.2        10.10.9.6  11/09/2010 16:15:54
                      002               OK               00            EgressRtr
    002       127.2.0.255       10.10.42.2        10.10.9.6  11/09/2010 16:15:54
                      002               OK               00            EgressRtr
```

```
    003       127.3.0.255        10.10.43.2         10.10.9.6  11/09/2010 16:15:54
                     002                OK                00            EgressRtr
    004       127.4.0.255        10.10.44.2         10.10.9.6  11/09/2010 16:15:54
                     002                OK                00            EgressRtr
    005       127.5.0.255        10.10.45.2         10.10.9.6  11/09/2010 16:15:54
                     002                OK                00            EgressRtr

    ldp-treetrace discovery state: Done
    ldp-treetrace discovery status: ' OK '
    Total number of discovered paths: 5
    Total number of probe-failed paths: 0
    Total number of failed traces: 0
*A:Dut-B#
 Total number of Hops: 2
```

Manual ldp tree-trace

The **tools dump ldp-treetrace manual-prefix** command displays entries discovered by a previously run ldp-treetrace manual test.

```
*A:Dut-B# tools dump ldp-treetrace manual-prefix 10.20.1.6/32
    Discovered Paths:
    ===================
    Id   PathDst          EgrNextHop        ReplyRtrAddr      DiscoveryTime
           DiscoveryTtl      ProbeState        ProbeTmOutCnt     RtnCode
    ===  ================  ================  ================  ===================
    001       127.1.0.255        10.10.41.2         10.10.9.6  11/09/2010 16:20:01
                     002                OK                00            EgressRtr
    002       127.2.0.255        10.10.42.2         10.10.9.6  11/09/2010 16:20:01
                     002                OK                00            EgressRtr
    003       127.3.0.255        10.10.43.2         10.10.9.6  11/09/2010 16:20:01
                     002                OK                00            EgressRtr
    004       127.4.0.255        10.10.44.2         10.10.9.6  11/09/2010 16:20:01
                     002                OK                00            EgressRtr
    005       127.5.0.255        10.10.45.2         10.10.9.6  11/09/2010 16:20:01
                     002                OK                00            EgressRtr


     ldp-treetrace discovery state: Done
     ldp-treetrace discovery status: ' OK '
     Total number of discovered paths: 5
     Total number of failed traces: 0
    *A:Dut-B#


*A:Dut-B# tools dump ldp-treetrace manual-prefix 10.20.1.6/32 path-destination 127.1.0.255
    FEC: 10.20.1.6/32  PathDst: 127.1.0.255
    ==================================================
    Protocol Legend: L - LDP, R - RSVP, U - Not Applicable

    HopId HopAddr          HopRouterId      TTL Label1  Label2  Label3  Label4  Label5
    ===== ================ ================ === ======= ======= ======= ======= =======
    006          10.10.9.6        10.20.1.6 002 131071L 000000U 000000U 000000U 000000U
    001         10.10.41.2        10.20.1.4 001 131069L 000000U 000000U 000000U 000000U

     Total number of Hops: 2

    *A:Dut-B#
```

persistence

Syntax persistence

Context tools>dump

Description This command enables the context to display persistence information for debugging purposes.

submgt

Syntax **submgt** [**record** *record-key*]

Context tools>dump>persistence

Description This command displays subscriber management persistence information.

summary

Syntax **summary**

Context tools>dump>persistence

Description The context to display persistence summary information for debugging purposes.

Sample Output

```
A:ALA-B# tools dump persistence summary
=====================================================================
Persistence Summary on Slot A
=====================================================================
Client             Location           Entries in use    Status
---------------------------------------------------------------------
xxxxxx           cf1:\l2_dhcp.pst    200               ACTIVE
---------------------------------------------------------------------
Persistence Summary on Slot B
=====================================================================
Client             Location           Entries in use    Status
---------------------------------------------------------------------
xxxxxx             cf1:\l2_dhcp.pst    200               ACTIVE
---------------------------------------------------------------------
A:ALA-B#
```

redundancy

Syntax **redundancy**

Context tools>dump

Description This command enables the context to dump tools for redundancy.

multi-chassis

Syntax **multi-chassis**

Context tools>dump>redundancy>multi-chassis

Description This command enables the context to dump tools for multi-chassis redundancy.

mc-endpoint

Syntax **mc-endpoint peer** *ip-address*

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis endpoint information.

Parameters **peer** *ip-address* **—** Specifies the peer's IP address.

mc-ring

Syntax **mc-ring**
mc-ring peer *ip-address* [**ring** *sync-tag*]

Context tools>dump>redundancy>multi-chassis

Description This command dumpsmulti-chassis ring information.

peer *ip-address* **—** Specifies the peer's IP address.

ring *sync-tag* **—** Specifies the ring's sync-tag created in the **config>redundancy>mc>peer>mcr>ring** context.

srrp-sync-database

Syntax **srrp-sync-database** [**instance** *instance-id*] [**peer** *ip-address*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps SRRP database information.

peer *ip-address* **—** Specifies the peer's IP address.

instance *instance-id* **—** Dumps information for the specified Subscriber Router Redundancy Protocol instance configured on this system.

Values 1 — 4294967295

sync-database

Syntax **sync-database** [**peer** *ip-address*] [**port** *port-id* **|** *lag-id*] [**sync-tag** *sync-tag*] [**application** *application*] [**detail**] [**type** *type*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps MCS database information.

peer *ip-address* **—** Specifies the peer's IP address.

port *port-id* | *lag-id* **—** Indicates the port or LAG ID to be synchronized with the multi-chassis peer.

Values *slot/mda/port* or lag-*lag-id*

sync-tag *sync-tag* **—** Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

application *application* **—** Specifies a particular multi-chassis peer synchronization protocol application.

Values

dhcp-server:	local dhcp server
igmp:	Internet group management protocol
igmp-snooping:	igmp-snooping
mc-ring:	multi-chassis ring
mld-snooping:	multicast listener discovery-snooping
srrp:	simple router redundancy protocol
sub-host-trk:	subscriber host tracking
sub-mgmt:	subscriber management

type *type* **—** Indicates the locally deleted or alarmed deleted entries in the MCS database per multi-chassis peer.

Values alarm-deleted, local-deleted

detail **—** Displays detailed information.

ppp

Syntax **ppp** *port-id*

Context tools>dump

Description This command displays PPP information for a port.

Default none

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

Values *port-id* *slot/mda/port*[.*channel*]

bundle-id:	bundle-*type-slot/mda.bundle-num*	
	bundle:	keyword
	type:	ppp
	bundle-num:	1 — 256
bpgrp-id:	bpgrp-*type-bpgrp-num*	
	bpgrp:	keyword
	type:	ppp

	bpgrp-num:	1 — 1280
aps-id:	aps-*group-id*[.channel]	
	aps:	keyword
	group-id:	1 — 64

Sample Output

```
*A:sr7# tools dump ppp aps-1.1.1.1
===============================================================================
Id              : aps-1.1.1.1        ppp unit       : 40
member of       : bpgrp-ppp-1
===============================================================================
looped back     : no                 dbgMask        : 0x0
-------------------------------------------------------------------------------
LCP
-------------------------------------------------------------------------------
phase           : NETWORK            state          : OPENED
passive         : off                silent         : off
restart         : on

mru             : 1500               mtu            : 1502
ack'd peer mru  : 1500
got local mrru  : 1524
local magic     : 0x0                peer magic     : 0x0

keepalive       : on                 echo num       : 2
echo timer      : on                 echos fail     : 3
echo intv       : 10                 echos pend     : 0

options      mru      asyncMap upap     chap     magic    pfc
we negotiate Yes       No       No       No       No      Yes
peer ack'd   Yes       No       No       No       No       No
we allow     Yes       No       No       No       No      Yes
we ack'd     Yes       No       No       No       No       No

options      acfc     lqr      mrru     shortSeq endPoint mlhdrfmt
we negotiate Yes       No      Yes       No      Yes       No
peer ack'd    No       No      Yes       No      Yes       No
we allow     Yes       No      Yes      Yes      Yes       No
we ack'd      No       No      Yes       No      Yes       No
...
===============================================================================
*A:sr7#
```

system-resources

Syntax **system-resources** *slot-number*

Context tools>dump

Description This command displays system resource information.

Default none

Parameters *slot-number* — Specify a specific slot to view system resources information.

Service Commands

service

Syntax **service**

Context tools>dump

Description Use this command to configure tools to display service dump information.

base-stats

Syntax **base-stats** [**clear**]

Context tools>dump>service

Description Use this command to display internal service statistics.

Default none

Parameters **clear —** Clears stats after reading.

iom-stats

Syntax **iom-stats** [**clear**]

Context tools>dump>service

Description Use this command to display IOM message statistics.

Default none

Parameters **clear —** Clears stats after reading.

l2pt-diags

Syntax **l2pt-diags**
l2pt-diags clear
l2pt-diags detail

Context tools>dump>service

Description Use this command to display L2pt diagnostics.

Default none

Parameters **clear —** Clears the diags after reading.

detail — Displays detailed information.

Sample Output

```
A:ALA-48>tools>dump>service# l2pt-diags
[ l2pt/bpdu error diagnostics ]
 Error Name       | Occurence   | Event log
 -----------------+-------------+----------------------------
[ l2pt/bpdu forwarding diagnostics ]

 Rx Frames   | Tx Frames   | Frame Type
 ------------+-------------+---------------------------------
A:ALA-48>tools>dump>service#

A:ALA-48>tools>dump>service# l2pt-diags detail
[ l2pt/bpdu error diagnostics ]
 Error Name       | Occurence   | Event log
 -----------------+-------------+----------------------------
[ l2pt/bpdu forwarding diagnostics ]

 Rx Frames   | Tx Frames   | Frame Type
 ------------+-------------+---------------------------------
[ l2pt/bpdu config diagnostics ]
 WARNING - service 700 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
 WARNING - service 800 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
 WARNING - service 9000 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
 WARNING - service 32806 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
 WARNING - service 90001 has l2pt termination enabled on all access points :
           consider translating further down the chain or turning it off.
A:ALA-48>tools>dump>service#
```

mc-endpoint

Syntax **mc-endpoint** *mc-ep-id*

Context tools>dump>service

Description Use this command to display multi-chassis endpoint information.

Parameters *mc-ep-id* — Specifies a multi-chassis endpoint ID.

Values 1 — 4294967295

Sample Output

```
*A:Dut-B#   tools dump service mc-endpoint 1
MC Endpoint Info
    mc-endpoint id              : 1
    endpoint                    : mcep-t1
    service                     : 1
```

```
    peer ref type                : peer-name
    peer                         : Dut-C
    mc sel logic                 : peer selected active
    selection master             : No
    retransmit pending           : No
    initial config sync          : Yes
    config sync                  : Yes
    peer not mcep                : No
    peer acked non-mcep          : No
    config mismatch              : No
    initial state rx             : Yes
    initial state sync           : Yes
    state sync                   : Yes
    can aggregate                : Yes
    sel peer active              : No
    peer sel active              : Yes
    passive mode active          : No
    own eligible force           : No
    own eligible double active   : Yes
    own eligible pw status bits  : 0
    own eligible precedence      : 2
    own eligible conf chg        : No
    own eligible revert wait     : No
    peer eligible force          : No
    peer eligible double active  : Yes
    peer eligible pw status bits : 0
    peer eligible precedence     : 3
    peer eligible conf chg       : No
    peer eligible revert wait    : No
*A:Dut-B# tools perform service id 1 endpoint mcep-t1 force-switchover 221:1
*A:Dut-B>show#
*A:Dut-B# show service  id  1 endpoint
===============================================================================
Service 1 endpoints
===============================================================================
Endpoint name                : mcep-t1
Description                  : (Not Specified)
Revert time                  : 0
Act Hold Delay               : 0
Ignore Standby Signaling     : false
Suppress Standby Signaling   : false
Block On Mesh Fail           : true
Multi-Chassis Endpoint       : 1
MC Endpoint Peer Addr        : 3.1.1.3
Psv Mode Active              : No
Tx Active                    : 221:1(forced)
Tx Active Up Time            : 0d 00:00:17
Revert Time Count Down       : N/A
Tx Active Change Count       : 6
Last Tx Active Change        : 02/14/2009 00:17:32
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 221:1 Prec:1                            Oper Status: Up
Spoke-sdp: 231:1 Prec:2                            Oper Status: Up
===============================================================================
*A:Dut-B#
```

radius-discovery

Syntax **radius-discovery** [**svc-id** *service-id*]

Context tools>dump>service

Description Use thie command to display RADIUS Discovery membership information.

Sample Output

```
A:ALA-48# tools dump service radius-discovery
--------------------------------------------------------------
Service Id 103  Vpn Id 103  UserName 901:103 (Vpn-Id)  PolicyName RAD_Disc for Service 103
Waiting for Session Timeout (Polling 60), Seconds in State 17
------------------------------------------------------------------------------------
      SdpId      Vcid  Deliver     Ip Addr        VcType       Mode      Split Horizon
-----------------------------------------------------------------------------------
          3       103    LDP    10. 20.  1.  3     Ether      Spoke
          4       103    LDP    10. 20.  1.  2     Ether      Spoke
 --------------------------------------------------------------
A:ALA-48#
```

vpls-fdb-stats

Syntax **vpls-fdb** [**clear**]

Context tools>dump>service

Description Use this command to display VPLS FDB statistics.

Default none

Parameters **clear —** Clears stats after reading.

vpls-mfib-stats

Syntax **vpls-mfib-stats** [**clear**]

Context tools>dump>service

Description Use this command to display VPLS MFIB statistics.

Default none

Parameters **clear —** Clears stats after reading.

Router Commands

router

Syntax **router** *router-instance*

Context tools>dump
tools>perform

Description This command enables tools for the router instance.

Default none

Parameters **router** *router-instance* **—** Specifies the router name or service ID.

Values *router-name*: Base , management
service-id: 1 — 2147483647

Default Base

dhcp

Syntax **dhcp**

Context tools>dump>router

Description This command enables the context to configure dump router tools for DHCP.

group-if-mapping

Syntax **group-if-mapping** [**clear**]

Context tools>dump>router>dhcp

Description This command dumps group interface mapping information stored in by the DHCP cache for the Routed CO model of operation.

group-if-stats

Syntax **group-if-stats** [**clear**]

Context tools>dump>router>dhcp

Description This command dumps group interface statistics information about the DHCP cache for the Routed CO model of operation.

lag

Syntax **lag**

Context tools>perform

Description This command configures tools to control LAG.

clear-force

Syntax **clear-force all-mc**
clear-force peer-mc *ip-address*
clear-force lag-id *lag-id* [**sub-group** *sub-group-id*]

Context tools>perform>lag

Description This command clears a forced status.

Parameters **all-mc —** Clears all multi-chassis LAG information.

lag-id ***lag-id*** **—** Specify an existing LAG id.

Values 1 — 200 (7750 SR-1: 1 — 64)

force

Syntax **force all-mc** {**active | standby**}
force peer-mc *peer-ip-address* {**active | standby**}
force lag-id *lag-id* [**sub-group** *sub-group-id*] {**active | standby**}

Context tools>perform>lag

Description This command forces an active or standy status.

Parameters **all-mc —** Clears all multi-chassis LAG information.

active — If **active** is selected, then all drives on the active CPM are forced.

standby — If **standby** is selected, then all drives on the standby CPM are forced.

all-mc — Clears all multi-chassis LAG information.

lag-id *lag-id* **—** Specify an existing LAG id.

Values 1 — 200 (7750 SR-1: 1 — 64)

log

Syntax **log**

Context tools>perform

Description Tools for event logging.

test-event

Syntax **test-event**

Context tools>perform>log

Description This command causes a test event to be generated. The test event is LOGGER event #2011 and maps to the tmnxEventSNMP trap in the TIMETRA-LOG-MIB.

ldp

Syntax **ldp**

Context tools>dump>router

Description This command enables dump tools for LDP.

Default none

interface

Syntax **interface** [*ip-int-name* | *ip-address*]

Context tools>dump>router>ldp

Description This command displays information for an LDP interface.

Default none

Parameters *ip-int-name* — Specifies the interface name.

ip-address — Specifies the IP address.

peer

Syntax **peer** *ip-address*

Context tools>dump>router>ldp

Description This command displays information for an LDP peer.

Default none

Parameters *ip-address* — Specifies the IP address.

fec

Syntax **fec prefix** [*ip-prefix*/*mask*]
fec vc-type {**ethernet | vlan**} **vc-id** *vc-id*

Context tools>dump>router>ldp

Description This command displays information for an LDP FEC.

Default none

Parameters *ip-prefix/mask* — Specifies the IP prefix and host bits.

Values host bits: must be 0
mask: 0 — 32

vc-type — Specifies the VC type signaled for the spoke or mesh binding to the far end of an SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the Dot1q value expected by the far-end provider equipment. A change of the binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

- Ethernet — The VC type value for Ethernet is 0x0005.
- VLAN — The VC type value for an Ethernet VLAN is 0x0004.

vc-id — Specifies the virtual circuit identifier.

Values 1 — 4294967295

instance

Syntax **instance**

Context tools>dump>router>ldp

Description This command displays information for an LDP instance.

memory-usage

Syntax **memory-usage**

Context tools>dump>router>ldp

Description This command displays memory usage information for LDP.

Default none

session

Syntax **session** [*ip-address* |*:label space*] [*connection* | *peer* | *adjacency*]

Context tools>dump>router>ldp

Description This command displays information for an LDP session.

Default none

Parameters *ip-address* — Specifies the IP address of the LDP peer.

label-space — Specifies the label space identifier that the router is advertising on the interface.

connection — Displays connection information.

peer — Displays peer information.

adjacency — Displays hello adjacency information.

sockets

Syntax **sockets**

Context tools>dump>router>ldp

Description This command displays information for all sockets being used by the LDP protocol.

Default none

timers

Syntax **timers**

Context tools>dump>router>ldp

Description This command displays timer information for LDP.

Default **none**

mpls

Syntax **mpls**

Context tools>dump>router

Description This command enables the context to display MPLS information.

Default none

ftn

Syntax **ftn**

Context tools>dump>router>mpls

Description This command displays FEC-to-NHLFE (FTN) dump information for MPLS. (NHLFE is the acronym for Next Hop Label Forwarding Entry.)

Default none

ilm

Syntax **ilm**

Context tools>dump>router>mpls

Description This command displays incoming label map (ILM) information for MPLS.

Default none

lspinfo

Syntax **lspinfo** [*lsp-name*] [**detail**]

Context tools>dump>router>mpls

Description This command displays label-switched path (LSP) information for MPLS.

Default none

Parameters *lsp-name* — Specifies the name that identifies the LSP. The LSP name can be up to 32 characters long and must be unique.

detail — Displays detailed information about the LSP.

memory-usage

Syntax **memory-usage**

Context tools>dump>router>mpls

Description This command displays memory usage information for MPLS.

Default none

te-lspinfo

Syntax **te-lspinfo** [**endpoint** *ip-address*] [**sender** *ip-address*] [**lspid** *lsp-id*] [**detail**] [**p2p** | ***p2p-tid*** *tunnel-id*]
te-lspinfo [**endpoint** *ip-address*] [**sender** *ip-address*] [**lspid** *lsp-id*] [**detail**] [**p2p** | **p2p-tid** *tunnel-id*]{ [**phops**] [**nhops**] [**s2l** *ip-address*] } }

Context tools>dump>router>mpls

Description This command displays TE LSP information for MPLS.

Default none

Sample Output

```
B:Dut-R# tools dump router mpls te-lspinfo
Key P2P: Session(10.10.3.2, 201, 3.3.3.3) Sender(3.3.3.3, 2) PHOP(10.10.3.1), Flags 0x0

Key P2P: Session(10.10.3.1, 1035, 4.4.4.4) Sender(4.4.4.4, 22) PHOP(10.10.11.2), Flags 0x0

Key P2MP: Session(0.0.0.0, 1, 4.4.4.4) Sender(4.4.4.4, 52226) PHOP(0.0.0.0) Flags 0x10
  S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
  S2L [2] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
  S2L [3] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

Key P2MP: Session(0.0.0.0, 2, 4.4.4.4) Sender(4.4.4.4, 51714) PHOP(0.0.0.0) Flags 0x10
  S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
  S2L [2] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
  S2L [3] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

Key P2MP: Session(0.0.0.0, 3, 4.4.4.4) Sender(4.4.4.4, 53250) PHOP(0.0.0.0) Flags 0x10

*A:Dut-T# tools dump router mpls te-lspinfo p2mp-tid 102 nhops

  Key P2MP: Session(0.0.0.0, 102, 4.4.4.4) Sender(4.4.4.4, 3074) PHOP(0.0.0.0) Flags 0x10
  -----------------------------------------------------------------------
        List of NEXT HOPS
  -----------------------------------------------------------------------

  NextHop [1] =>
  Key: Nhop - isFrr 0, outIf 0, NextHop 0.0.0.0 label - 128843  global Instance 0 is Leaf
node
        -----------------------------------------------------------------------
        Primary NHLFE => outLabel - 0 and NextHop - 0.0.0.0, outIf 0 (0)
                Port(NONE) NhIdx 0, ProtNhIdx 0, NumS2L 1
                ProtectInstance - 0, ProtectGroup 0
        POP
        No Backup NHLFEs for this Ltn entry
  Mid List :    3428 numS2Ls - 1 (Primary MID),

  NextHop [2] =>
  Key: Nhop - isFrr 0, outIf 3, NextHop 10.10.13.2 label - 128806  global Instance -48747
        -----------------------------------------------------------------------
        Primary NHLFE => outLabel - 128806 and NextHop - 10.10.13.2, outIf 3 (126)
                Port(9/1/1) NhIdx 4322, ProtNhIdx 2275, NumS2L 1
                ProtectInstance - 1, ProtectGroup 126
        SWAP
        Backup NHLFE => outLabel - 130223 and NextHop - 10.10.3.2, outIf 5 (124)
```

```
                Port(9/2/3) outPushLabel 128806, NhIdx 5469, ProtNhIdx 0, NumS2L 1
Mid List :      3428 numS2Ls - 1 (Primary MID),

NextHop [3] =>
Key: Nhop - isFrr 0, outIf 4, NextHop 10.10.2.2 label - 128836  global Instance -48974
      ----------------------------------------------------------------------
      Primary NHLFE => outLabel - 128836 and NextHop - 10.10.2.2, outIf 4 (125)
                Port(lag-1) NhIdx 4292, ProtNhIdx 2245, NumS2L 2
                ProtectInstance - 1, ProtectGroup 125
      SWAP
      Backup NHLFE => outLabel - 130223 and NextHop - 10.10.3.2, outIf 5 (124)
                Port(9/2/3) outPushLabel 128836, NhIdx 5659, ProtNhIdx 0, NumS2L 2
Mid List :      3428 numS2Ls - 1 (Primary MID),   3471 numS2Ls - 1 (Backup MID),

S2L [1] Key: endPoint to 2.2.2.2 subGroupId - 1 subGroupOrigId - 4.4.4.4
S2L [2] Key: endPoint to 3.3.3.3 subGroupId - 2 subGroupOrigId - 4.4.4.4
S2L [3] Key: endPoint to 10.10.2.2 subGroupId - 3 subGroupOrigId - 4.4.4.4
S2L [4] Key: endPoint to 10.10.13.2 subGroupId - 4 subGroupOrigId - 4.4.4.4

Total TeLspInfo Count   : 1
```

ospf

Syntax **ospf** [*ospf-instance*]

Context tools>dump>router

Description This command enables the context to display tools information for OSPF.

Default none

Parameters **ospf-instance —** OSPF instance.

Values 1 — 4294967295

ospf3

Syntax **ospf3**

Context tools>dump>router

Description This command enables the context to display tools information for OSPF3.

Default none

abr

Syntax **abr** [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays area border router (ABR) information for OSPF.

Default none

Parameters **detail —** Displays detailed information about the ABR.

asbr

Syntax **asbr** [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays autonoumous system border router (ASBR) information for OSPF.

Default none

Parameters **detail —** Displays detailed information about the ASBR.

bad-packet

Syntax **bad-packet** [*interface-name*]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays information about bad packets for OSPF.

Default none

Parameters *interface-name* — Display only the bad packets identified by this interface name.

leaked-routes

Syntax **leaked-routes** [**summary | detail**}

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays information about leaked routes for OSPF.

Default summary

Parameters **summary —** Display a summary of information about leaked routes for OSPF.

detail — Display detailed information about leaked routes for OSPF.

memory-usage

Syntax **memory-usage** [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays memory usage information for OSPF.

Default none

Parameters **detail —** Displays detailed information about memory usage for OSPF.

request-list

Syntax **request-list** [**neighbor** *ip-address*] [**detail**]
request-list virtual-neighbor *ip-address* **area-id** *area-id* [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays request list information for OSPF.

Default none

Parameters **neighbor** *ip-address* — Display neighbor information only for neighbor identified by the IP address.

detail — Displays detailed information about the neighbor.

virtual-neighbor *ip-address* — Displays information about the virtual neighbor identified by the IP address.

area-id *area-id* — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

retransmission-list

Syntax **retransmission-list** [**neighbor** *ip-address*] [**detail**]
retransmission-list virtual-neighbor *ip-address* **area-id** *area-id* [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays dump retransmission list information for OSPF.

Default none

Parameters **neighbor** *ip-address* — Display neighbor information only for neighbor identified by the IP address.

detail — Displays detailed information about the neighbor.

virtual-neighbor *ip-address* — Displays information about the virtual neighbor identified by the IP address.

area-id *area-id* — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

route-summary

Syntax **route-summary**

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays dump route summary information for OSPF.

Default none

route-table

Syntax **route-table** [**type**] [**detail**]

Context tools>dump>router>ospf
tools>dump>router>ospf3

Description This command displays dump information about routes learned through OSPF.

Default none

Parameters **type —** Specify the type of route table to display information.

Values intra-area, inter-area, external-1, external-2, nssa-1, nssa-2

detail — Displays detailed information about learned routes.

pim

Syntax **pim**

Context tools>dump>router

Description This command enables the context to display PIM information.

iom-failures

Syntax **iom-failures** [**detail**]

Context tools>dump>router>pim

Description This command displays information about failures in programming IOMs.

Parameters *detail* — Displays detailed information about IOM failures.

rsvp

Syntax rsvp

Context tools>dump>router

Description This command enables the context to display RSVP information.

Default none

psb

Syntax **psb** [**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** *tunnel-id*] [**lspid** *lsp-id*]

Context tools>dump>router>rsvp

Description This command displays path state block (PSB) information for RSVP.

When a PATH message arrives at an LSR, the LSR stores the label request in the local PSB for the LSP. If a label range is specified, the label allocation process must assign a label from that range.

The PSB contains the IP address of the previous hop, the session, the sender, and the TSPEC. This information is used to route the corresponding RESV message back to LSR 1.

Default none

Parameters **endpoint** *endpoint-address* — Specifies the IP address of the last hop.

sender *sender-address* — Specifies the IP address of the sender.

tunnelid *tunnel-id* — Specifies the SDP ID.

Values 0 — 4294967295

lspid *lsp-id* — Specifies the label switched path that is signaled for this entry.

Values 1 — 65535

rsb

Syntax **rsb** [**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** *tunnel-id*] [**lspid** *lsp-id*]

Context tools>dump>router>rsvp

Description This command displays RSVP Reservation State Block (RSB) information.

Default none

Parameters **endpoint** *endpoint-address* — Specifies the IP address of the last hop.

sender *sender-address* — Specifies the IP address of the sender.

tunnelid *tunnel-id* — Specifies the SDP ID.

Values 0 — 4294967295

lspid *lsp-id* — Specifies the label switched path that is signaled for this entry.

Values 1 — 65535

tcsb

Syntax **tcsb** [**endpoint** *endpoint-address*] [**sender** *sender-address*] [**tunnelid** *tunnel-id*] [**lspid** *lsp-id*]

Context tools>dump>router>rsvp

Description This command displays RSVP traffic control state block (TCSB) information.

Default none

Parameters **endpoint** *endpoint-address* — Specifies the IP address of the egress node for the tunnel supporting this session.

sender *sender-address* — Specifies the IP address of the sender node for the tunnel supporting this session. It is derived from the source address of the associated MPLS LSP definition.

tunnelid *tunnel-id* — Specifies the IP address of the ingress node of the tunnel supporting this RSVP session.

Values 0 — 4294967295

lspid *lsp-id* — Specifies the label switched path that is signaled for this entry.

Values 1 — 65535

static-route

Syntax **static-route ldp-sync-status**

Context tools>dump>router

Description This command displays the sync status of LDP interfaces that static-route keeps track of.

web-rd

Syntax **web-rd**

Context tools>dump>router

Description This command enables the context to display tools for web redirection.

http-client

Syntax **http-client** [*ip-prefix/mask*]

Context tools>dump>router>web-rd

Description This command displays the HTTP client hash table.

Parameters *ip-prefix/mask* — Specifies the IP prefix and host bits.

Values		
	host bits:	must be 0
	mask:	0 — 32

Performance Tools

perform

Syntax perform

Context tools

Description This command enables the context to enable tools to perform specific tasks.

Default none

cron

Syntax **cron**

Context tools>perform

Description This command enables the context to perform CRON (scheduling) control operations.

Default none

action

Syntax **action**

Context tools>perform>cron

Description This command enables the context to stop the execution of a script started by CRON action. See the **stop** command.

stop

Syntax **stop** [*action-name*] [**owner** *action-owner*] [**all**]

Context tools>perform>cron>action

Description This command stops execution of a script started by CRON action.

Parameters *action-name* — Specifies the action name.

Values Maximum 32 characters.

owner *action-owner* **—** Specifies the owner name.

Default TiMOS CLI

all **—** Specifies to stop all CRON scripts.

tod

Syntax **tod**

Context tools>perform>cron

Description This command enables the context for tools for controlling time-of-day actions.

Default none

re-evaluate

Syntax **re-evaluate**

Context tools>perform>cron>tod

Description This command enables the context to re-evaluate the time-of-day state.

Default none

customer

Syntax **customer** *customer-id* [**site** *customer-site-name*]

Context tools>perform>cron>tod>re-eval

Description This command re-evaluates the time-of-day state of a multi-service site.

Parameters *customer-id* — Specify an existing customer ID.

Values 1 — 2147483647

site *customer-site-name* **—** Specify an existing customer site name.

filter

Syntax **filter** *filter-type* [*filter-id*]

Context tools>perform>cron>tod>re-eval

Description This command re-evaluates the time-of-day state of a filter entry.

Parameters *filter-type* — Specify the filter type.

Values ip-filter, ipv6-filter, mac-filter

filter-id — Specify an existing filter ID.

Values 1 — 65535

service

Syntax **service id** *service-id* [**sap** *sap-id*]

Context tools>perform>cron>tod>re-eval

Description This command re-evaluates the time-of-day state of a SAP.

Parameters **id** *service-id* — Specify the an existing service ID.

Values 1 — 2147483647

sap ***sap-id*** **—** Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 355 for CLI command syntax.

tod-suite

Syntax **tod-suite** *tod-suite-name*

Context tools>perform>cron>tod>re-eval

Description This command re-evaluates the time-of-day state for the objects referring to a tod-suite.

Parameters *tod-suite-name* — Specify an existing TOD nfame.

aps

Syntax **aps**

Context tools>perform

Description This command enables the context to perform Automated Protection Switching (APS) operations.

clear

Syntax **clear** *aps-id* {**protect** | **working**}

Context tools>perform>aps

Description This command removes all Automated Protection Switching (APS) operational commands.

Parameters *aps-id* — This option clears a specific APS on un-bundled SONET/SDH ports.

protect — This command clears a physical port that is acting as the protection circuit for the APS group.

working — This command clears a physical port that is acting as the working circuit for this APS group.

exercise

Syntax **exercise** *aps-id* {**protect | working**}

Context tools>perform

Description This command performs an exercise request on the protection or working circuit.

Parameters *aps-id* — This option clears a specific APS on un-bundled SONET/SDH ports.

protect — This command performs an exercise request on the port that is acting as the protection circuit for the APS group.

working — This command performs an exercise request on the port that is acting as the working circuit for this APS group.

force

Syntax **force** *aps-id* {**protect | working**}

Context tools>perform

Description This command forces a switch to either the protect or working circuit

Parameters *aps-id* — This option clears a specific APS on un-bundled SONET/SDH ports.

protect — This command clears a physical port that is acting as the protection circuit for the APS group.

working — This command clears a physical port that is acting as the working circuit for this APS group.

lockout

Syntax **lockout** *aps-id*

Context tools>perform

Description This command locks out the protection circuit.

Parameters *aps-id* — Automated Protection Switching ID

Values 1 — 64

request

Syntax **request** *aps-id* {**protect | working**}

Context tools>perform

Description This command requests a manual switch to protection or working circuit.

Parameters *aps-id* — This option clears a specific APS on un-bundled SONET/SDH ports.

protect — This command requests a manual switch to a port that is acting as the protection circuit for the APS group.

working — This command requests a manual switch to a port that is acting as the working circuit for this APS group.

consistency

Syntax **consistency**

Context tools>perform>router

Description This command performs route table manager (RTM) consistency checks.

Default none

ldp-sync-exit

Syntax [**no**] **ldp-sync-exit**

Context tools>perform>router>isis
tools>perform>router>ospf

Description This command restores the actual cost of an interface at any time. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces which have the IGP-LDP synchronization enabled if the currently advertised cost is different.

isis

Syntax **isis**

Context tools>perform>router

Description This command enables the context to configure tools to perform certain ISIS tasks.

run-manual-spf

Syntax **run-manual-spf**

Context tools>perform>router>isis

Description This command runs the Shortest Path First (SPF) algorithm.

mpls

Syntax **mpls**

Context tools>perform>router

Description This command enables the context to perform specific MPLS tasks.

Default none

adjust-autobandwidth

Syntax **adjust-autobandwidth** [**lsp** *lsp-name* [**force** [**bandwidth** *mbps*]]]

Context tools>perform>router>mpls

Description This command initiates an immediate auto-bandwidth adjustment attempt for either one specific LSP or all active LSPs. If an LSP is not specified then the system assumes the command applies to all LSPs.

The adjust-count, maximum average data rate and overflow count are not reset by the manual auto-bandwidth command, whether or not the bandwidth adjustment succeeds or fails.

Parameters **lsp** *lsp-name* **—** Specifies the name that identifies the LSP. The LSP name can be up to 32 characters longand must be unique.

force — The optional force parameter, which is available only when an LSP is referenced, determines whether adjust-up and adjust-down threshold checks are applied. If force is not specified then the maximum average data rate must differ from the current reservation by more than the adjust-up or adjust-down thresholds, otherwise no bandwidth adjustment occurs. If the force option is specified then, bandwidth adjustment ignores the configured thresholds.

bandwidth *mbps* **—** If a bandwidth is specified as part of the force option then the bandwidth of the LSP is changed to this specific value, otherwise the bandwidth is changed to the maximum average data rate that has been measured by the system in the current adjust interval.

cspf

Syntax **cspf to** *ip-addr* [**from** *ip-addr*] [**bandwidth** *bandwidth*] [**include-bitmap** *bitmap*] [**exclude-bitmap** *bitmap*] [**hop-limit** *limit*] [**exclude-address** *excl-addr* [*excl-addr*...(up to 8 max)]] [**use-te-metric**] [**strict-srlg**] [**srlg-group** *grp-id*...(up to 8 max)] [**exclude-node** *excl-node-id* [*excl-node-id* ..(up to 8 max)]] [**skip-interface** *interface-name*] [**ds-class-type** *class-type*] [**cspf-reqtype** *req-type*] [**least-fill-min-thd** *thd*] [**setup-priority** *val*] [**hold-priority** *val*]

Context tools>perform>router>mpls

Description This command computes a CSPF path with specified user constraints.

Default none

Parameters **to** *ip-addr* **—** Specify the destination IP address.

from *ip-addr* **—** Specify the originating IP address.

bandwidth *bandwidth* **—** Specifies the amount of bandwidth in mega-bits per second (Mbps) to be reserved.

include-bitmap *bitmap* **—** Specifies to include a bit-map that specifies a list of admin groups that should be included during setup.

exclude-bitmap *bitmap* **—** Specifies to exclude a bit-map that specifies a list of admin groups that should be included during setup.

hop-limit *limit* **—** Specifies the total number of hops a detour LSP can take before merging back onto the main LSP path.

exclude-address *ip-addr* **—** Specifies an IP address to exclude from the operation.

use-te-metric — Specifies whether the TE metric would be used for the purpose of the LSP path computation by CSPF.

skip-interface *interface-name* **—** Specifies a local interface name, instead of the interface address, to be excluded from the CSPF computation.

ds-class-type *class-type* **—** Specifies the class type.

Values 0 — 7

cspf-reqtype *req-typ* **—** Specifies the CSPF request type.

Values all — Specifies all ECMP paths.
random — Specifies random ECMP paths.
least-fill — Specifies minimum fill path.

resignal

Syntax **resignal lsp** *lsp-name* **path** *path-name* **delay** *minutes*
resignal {**p2mp-lsp** *p2mp-lsp-name* **p2mp-instance** *p2mp-instance-name* **|** **p2mp-delay** *p2mp-minutes*}

Context tools>perform>router>mpls

Description Use this command to resignal a specific LSP path.

Default none

Parameters **lsp** *lsp-name* **—** Specifies the name that identifies the LSP. The LSP name can be up to 32 characters longand must be unique.

path *path-name* **—** Specifies the name for the LSP path up, to 32 characters in length.

delay *minutes* **—** Specifies the resignal delay in minutes.

Values 0 — 30

p2mp-lsp *p2mp-lsp-name* **—** Specifies an existing point-to-multipoint LSP name.

p2mp-instance *p2mp-instance-name* **—** Specifies a name that identifies the P2MP LSP instance

p2mp-delay *p2mp-minutes* **—** Specifies the delay time, in minutes.

Values 0 — 60

trap-suppress

Syntax **trap-suppress** [*number-of-traps*] [*time-interval*]

Context tools>perform>router>mpls

Description This command modifies thresholds for trap suppression.

Default none

Parameters *number-of-traps* — Specify the number of traps in multiples of 100. An error messages is generated if an invalid value is entered.

Values 100 to 1000

time-interval — Specify the timer interval in seconds.

Values 1 — 300

ospf

Syntax ospf

Context tools>perform>router

Description This command enables the context to perform specific OSPF tasks.

Default none

ospf3

Syntax ospf3

Context tools>perform>router

Description This command enables the context to perform specific OSPF3 tasks.

Default none

refresh-lsas

Syntax **refresh-lsas** [*lsa-type*] [*area-id*]

Context tools>perform>router>ospf
tools>perform>router>ospf3

Description This command refreshes LSAs for OSPF.

Default none

Parameters *lsa-type* — Specify the LSA type using allow keywords.

Values router, network, summary, asbr, extern, nssa, opaque

area-id — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

Values 0 — 4294967295

run-manual-spf

Syntax **run-manual-spf** *externals-only*

Context tools>perform>router>ospf
tools>perform>router>ospf3

Description This command runs the Shortest Path First (SPF) algorithm.

Default none

Parameters **externals-only —** Specify the route preference for OSPF external routes.

security

Syntax **security**

Context tools>perform

Description This command provides tools for testing security.

authentication-server-check

Syntax **authentication-server-check** *server-address ip-address* [**port** *port*] **user-name** *DHCP client user name* **password** *password* **secret** *key* [**source-address** *ip-address*] [**timeout** *seconds*] [**router** *router-instance*]

Context tools>perform>security

Description This command checks connection to the RADIUS server.

Parameters **router** *router-instance* **—** Specifies the router name or service ID.

Values *router-name*: Base , management
service-id: 1 — 2147483647

Default Base

service

Syntax **services**

Context tools>perform

Description This command enables the context to configure tools for services.

egress-multicast-group

Syntax **egress-multicast-group** *group-name*

Context tools>perform>service

Description This command enables the context to configure tools for egress multicast groups.

Parameters *group-name* — Specify an existing group name.

force-optimize

Syntax **force-optimize**

Context tools>perform>service>egress-multicast-group

Description This command optimizes the chain length.

eval-pw-template

Syntax **eval-pw-template** *policy-id* [**allow-service-impact**]

Context tools>perform>service>egress-multicast-group
tools>perform>service>id

Description This command re-evaluates the pseudowire template policy.

Parameters *policy-id* — Specifies the pseudowire template policy.

id

Syntax **id** *service-id*

Context tools>perform>service

Description This command enables the context to configure tools for a specific service.

Parameters *service-id* — Specify an existing service ID.

Values 1 — 2147483647

endpoint

Syntax **endpoint** *endpoint-name*

Context tools>perform>service>id

Description This command enables the context to configure tools for a specific VLL service endpoint.

Parameters *endpoint-name* — Specify an existing VLL service endpoint name.

force-switchover

Syntax **force-switchover** *sdp-id:vc-id*
no force-switchover

Context tools>perform>service>id

Description This command forces a switch of the active spoke SDP for the specified service.

Parameters *sdp-id:vc-id* — Specify an existing spoke SDP for the service.

Sample Output

```
A:Dut-B# tools perform service id 1 endpoint mcep-t1 force-switchover 221:1
*A:Dut-B# show service  id  1 endpoint
===============================================================================
Service 1 endpoints
===============================================================================
Endpoint name                : mcep-t1
Description                  : (Not Specified)
Revert time                  : 0
Act Hold Delay               : 0
Ignore Standby Signaling     : false
Suppress Standby Signaling   : false
Block On Mesh Fail           : true
Multi-Chassis Endpoint       : 1
MC Endpoint Peer Addr        : 3.1.1.3
Psv Mode Active              : No
Tx Active                    : 221:1(forced)
Tx Active Up Time            : 0d 00:00:17
Revert Time Count Down       : N/A
Tx Active Change Count       : 6
Last Tx Active Change        : 02/14/2009 00:17:32
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 221:1 Prec:1                             Oper Status: Up
Spoke-sdp: 231:1 Prec:2                             Oper Status: Up
================================================================================
*A:Dut-B#
```

mcac

Syntax **mcac sap** *sap-id* **recalc policy** *policy-name* [**bundle** *bundle-name*]
mcac sdp *sdp-id:vc-id* **recalc policy** *policy-name* [**bundle** *bundle-name*]

Context tools>perform>service>id

Description This command enables too for a multicast CAC.

Parameters **sap** *sap-id* — Specifies the SAP ID.

recalc — keyword

policy *policy-name* — Specifies the policy name.

bundle *bundle-**name*** — Specifies the bundle name.

subscriber-mgmt

Syntax **subscriber-mgmt**

Context tools>perform

Description This command enables tools to control subscriber management.

edit-lease-state

Syntax **edit-lease-state sap** *sap-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*]
edit-lease-state svc-id *service-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*]

Context tools>perform>subscr-mgmt

Parameters **sap** ***sap-id*** — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 355 for CLI command syntax.

ip ***ip-address*** — Modifies lease state information for the specified IP address.

subscriber ***sub-ident-string*** — Modifies lease state information for the specified subscriber ID.

sub-profile-string ***sub-profile-string*** — Modifies lease state information for the specified subscriber profile.

sla-profile-string ***sla-profile-string*** — Modifies lease state information for the SLA profile.

svc-id ***service-id*** — Modifies lease state information for the specified service ID.

Values 1 — 2147483647

eval-lease-state

Syntax **eval-lease-state** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ip** *ip-address*]

Context tools>perform>subscr-mgmt

Description This command evaluates lease state information.

Parameters **svc-id** ***service-id*** **—** Evaluates lease state information for the specified service.

Values 1 — 2147483647

sap ***sap-id*** **—** Evaluates lease state information for the specified SAP. See Common CLI Command Descriptions on page 355 for CLI command syntax.

subscriber ***sub-ident-string*** **—** Evaluates lease state information for the specified subscriber identification string.

ip ***ip-address*** **—** Evaluates lease state information for the specified IP address.

forcerenew

Syntax **forcerenew svc-id** *service-id* {**ip** *ip-address*[/*mask*] **| mac** *ieee-address*}
forcerenew {**interface** *interface-name* **| sap** *sap-id* **| sdp** *sdp-id:vc-id*} [**ip** *ip-address*[/*mask*] **|mac** *ieee-address*]

Context tools>perform>subscr-mgmt

Description This command forces the renewal of lease state.

Parameters **svc-id** ***service-id*** **—** Forces renewal of the lease state for the specified service.

Values 1 — 2147483647

sap ***sap-id*** **—** Forces renewal of the lease state for the specified SAP. See Common CLI Command Descriptions on page 355 for CLI command syntax.

ip ***ip-address*** **—** Forces renewal of the lease state for the specified IP address.

mac *ieee-address* **—** Forces renewal of the lease state for the specified MAC address.

interface *interface-name* **—** Forces renewal of the lease state for the specified interface name.

re-ident-sub

Syntax **re-ident-sub** *old-sub-ident-string* **to** *new-sub-ident-string*

Context tools>perform>subscr-mgmt

Description This command renames a subscriber identification string.

Parameters *old-sub-ident-string* — Specifies the existing subscriber identification string to be renamed.

new-sub-ident-string — Specifies the new subscriber identification string name.

remap-lease-state

Syntax **remap-lease-state old-mac** *ieee-address* **mac** *ieee-address*
remap-lease-state sap *sap-id* [mac *ieee-address*]

Context tools>perform>subscr-mgmt

Description This command allows the remapping of all existing hosts if network card on CMTS/WAC side is changed is required.

When this command is executed, the following restrictions apply

- When **sap** is taken, all leases associated with the SAP are re-written.
  - → For a SAP with a configured MAC in "lease-populate" command, this MAC will be taken.
  - → For a SAP without a configured MAC the MAC from tools command will be taken.
  - → For a SAP without a configured MAC and no MAC in tools command no action will be perform.
- When using the **old-mac** option, providing a new MAC *ieee-address* is mandatory.

This command is applicable only when dealing with DHCP lease states which were instantiated using l2header mode of DHCP operation.

Parameters **old-mac** *ieee-address*

old-mac *ieee-address* **—** specifies the old MAC address to remap.

mac *ieee-address* **—** Specifies that the provisioned MAC address will be used in the anti-spoofing entries for this SAP when l2-header is enabled. The parameter may be changed mid-session. Existing sessions will not be re-programmed unless a **tools perform** command is issued for the lease.

sap ***sap-id*** **—** Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 355 for CLI command syntax.

When configured, the SAP parameter will remap all MAC addresses of DHCP lease states on the specified SAP. When no optional MAC parameter is specified, the **sap** *sap-id* command remaps all MAC addresses of lease states towards the MAC address specified in the l2-header configuration.

Common CLI Command Descriptions

In This Chapter

This chapter provides CLI syntax and command descriptions for SAP and port commands.

Topics in this chapter include:

Common Service Commands

sap

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

Description This command specifies the physical port identifier portion of the SAP definition.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

The *sap-id* can be configured in one of the following formats:

Type	Syntax	Example
port-id	*slot*/*mda*/*port*[.*channel*]	1/1/5
null	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]	*port-id*: 1/1/3 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*: lag-3 *aps-id*: aps-1
dot1q	[*port-id* | *bundle-id*| *bpgrp-id* | *lag-id* | *aps-id*]:qtag1	*port-id*:qtag1: 1/1/3:100 *bundle-id*: bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1:lag-3:102 *aps-id*:qtag1: aps-1:27
qinq	[*port-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 1/1/3:100.10 *bpgrp-id*: bpgrp-ima-1 *lag-id*:qtag1.qtag2: lag-10:
atm	[*port-id* | *aps-id* | *bundle-id* | *bpgrp-id*][:vpi/vci |vpi |vpi1.vpi2]	port-id: 1/1/1 aps-id: aps-1 *bundle-id*: bundle-ima-1/1.1 bundle-ppp-1/1.1 *bpgrp-id*: bpgrp-ima-1 vpi/vci: 16/26 vpi: 16 vpi1.vpi2: 16.200
frame-relay	[*port-id* | *aps-id*]:*dlci*	*port-id*: 1/1/1:100 *bundle-id*: bundle-fr-3/1.1:100 *aps-id*: aps-1 *dlci*: 16
cisco-hdlc	*slot/mda/port.channel*	*port-id*: 1/1/3.1

7750 SR:

Values: *sap-id*	null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
	dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
	qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
	atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*| *vpi1.vpi2*]	
	frame	[*port-id* | *aps-id*]:*dlci*	
	cisco-hdlc	*slot/mda/port.channel*	
	cem	*slot/mda/port.channel*	
	ima-grp	[*bundle-id*[:vpi/vci|vpi|*vpi1.vpi2*]	
	port-id	*slot/mda/port*[.*channel*]	
	bundle-id	bundle-*type-slot/mda.bundle-num*	
		bundle	keyword
		type	ima, fr, ppp
		bundle-num	1 — 336
	bpgrp-id	bpgrp-*type-bpgrp-num*	
		bpgrp	keyword
		type	ima, ppp
		bpgrp-num	1 — 2000
	aps-id	aps-*group-id*[.*channel*]	
		aps	keyword
		group-id	1 — 64
	ccag-id	ccag-*id*.*path-id*[*cc-type*]:*cc-id*	
		ccag	keyword
		id	1 — 8
		path-id	a, b
		cc-type	.sap-net, .net-sap
		cc-id	0 — 4094
	eth-tunnel	eth-tunnel-*id*[:*eth-tun-sap-id*]	
		id	1— 1024
		eth-tun-sap-id	0 — 4094
	lag-id	lag-id	
		lag	keyword
		id	1 — 200
	qtag1	0 — 4094	
	qtag2	*, 0 — 4094	
	vpi	NNI: 0 — 4095	
		UNI: 0 — 255	
	vci	1, 2, 5 — 65535	
	dlci	16 — 1022	
	ipsec-id	ipsec-*id*.[private | public]:*tag*	
		ipsec	keyword
		id	1 — 4
		tag	0 — 4094

7710 SR:

Values: *sap-id*:	null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
	dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
	qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
	atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*| *vpi1.vpi2*]	
	frame	[*port-id* | *aps-id*]:*dlci*	
	cisco-hdlc	*slot/mda/port.channel*	
	cem	*slot/mda/port.channel*	
	ima-grp	[*bundle-id*[:vpi/vci|vpi|*vpi1.vpi2*]	
	port-id	*slot/mda/port*[.*channel*]	
	bundle-id	bundle-*type-slot/mda.bundle-num*	
		bundle	keyword
		type	ima, ppp
		bundle-num	1 — 256
	bpgrp-id	bpgrp-*type-bpgrp-num*	
		bpgrp	keyword
		type	ima, ppp
		bpgrp-num	1 — 1280
	aps-id	aps-*group-id*[.*channel*]	
		aps	keyword
		group-id	1 — 16
	lag-id	lag-*id*	
		lag	keyword
		id	1 — 64
	qtag1	0 — 4094	
	qtag2	*, 0 — 4094	
	vpi	NNI: 0 — 4095	
		UNI: 0 — 255	
	vci	1, 2, 5 — 65535	
	dlci	16 — 1022	

7450 ESS:

Values: *sap-id*	null	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]	
	dot1q	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id* | *aps-id*]:*qtag1*	
	qinq	[*port-id* | *bundle-id* | *bpgrp-id* | *lag-id*]:*qtag1.qtag2*	
	atm	[*port-id* | *aps-id*][:*vpi/vci*|*vpi*| *vpi1.vpi2*]	
	frame	[*port-id* | *aps-id*]:*dlci*	
	cisco-hdlc	*slot/mda/port.channel*	
	ima-grp	[*bundle-id*[:vpi/vci|vpi|*vpi1.vpi2*]	
	port-id	*slot/mda/port*[.*channel*]	
	bundle-id	bundle-*type-slot/mda.bundle-num*	
		bundle	keyword
		type	ima, fr, ppp
		bundle-num	1 — 336
	bpgrp-id	bpgrp-*type-bpgrp-num*	
		bpgrp	keyword
		type	ima, ppp
		bpgrp-num	1 — 2000
	aps-id	aps-*group-id*[.*channel*]	
		aps	keyword
		group-id	1 — 64
	ccag-id	ccag-*id.path-id*[*cc-type*]:*cc-id*	
		ccag	keyword
		id	1 — 8
		path-id	a, b
		cc-type	.sap-net, .net-sap
		cc-id	0 — 4094
	eth-tunnel	eth-tunnel-*id*[:*eth-tun-sap-id*]	
		id	1— 1024
		eth-tun-sap-id	0 — 4094
	lag-id	lag-id	
		lag	keyword
		id	1 — 200
	qtag1	0 — 4094	
	qtag2	*, 0 — 4094	
	vpi	NNI: 0 — 4095	
		UNI: 0 — 255	
	vci	1, 2, 5 — 65535	
	dlci	16 — 1022	

port

Syntax **port** *port-id*

Description This command specifies a port identifier.

Parameters *port-id* — The *port-id* can be configured in one of the following formats.

Values

port-id	slot/mda/port[.channel]		
	bundle-id	bundle-*type-slot/mda.bundle-num*	
		bundle	keyword
		type	ima|ppp
		bundle-num	1 — 256
	bpgrp-id	bpgrp-*type-bpgrp-num*	
		bpgrp	keyword
		type	ima, ppp
		bpgrp-num	1 — 256
	aps-id	aps-*group-id*[.*channel*]	
		aps	keyword
		group-id	1 — 64
	ccag-id	ccag-*id.*<*path-id*>[*cc-type*]	
		ccag	keyword
		id	1 — 8
		path-id	a, b
		cc-type	[.sap-net|.net-sap]
	lag-id	lag-*id*	
		lag	keyword
		id	1— 200

Standards and Protocol Support

Standards Compliance

IEEE 802.1ab-REV/D3 Station and Media Access Control Connectivity Discovery
IEEE 802.1d Bridging
IEEE 802.1p/Q VLAN Tagging
IEEE 802.1s Multiple Spanning Tree
IEEE 802.1w Rapid Spanning Tree Protocol
IEEE 802.1x Port Based Network Access Control
IEEE 802.1ad Provider Bridges
IEEE 802.1ah Provider Backbone Bridges
IEEE 802.1ag Service Layer OAM
IEEE 802.3ah Ethernet in the First Mile
IEEE 802.1ak Multiple MAC Registration Protocol
IEEE 802.3 10BaseT
IEEE 802.3ad Link Aggregation
IEEE 802.3ae 10Gbps Ethernet
IEEE 802.3ah Ethernet OAM
IEEE 802.3u 100BaseTX
IEEE 802.3x Flow Control
IEEE 802.3z 1000BaseSX/LX
ITU-T Y.1731 OAM functions and mechanisms for Ethernet based networks
ITU-T G.8031 Ethernet linear protection switching
ITU-T G.8032 Ethernet Ring Protection Switching (version 2)

Protocol Support

OSPF

RFC 1765 OSPF Database Overflow
RFC 2328 OSPF Version 2
RFC 2370 Opaque LSA Support
RFC 2740 OSPF for IPv6 (OSPFv3) draft-ietf-ospf-ospfv3-update-14.txt
RFC 3101 OSPF NSSA Option
RFC 3137 OSPF Stub Router Advertisement
RFC 3623 Graceful OSPF Restart – GR helper
RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2
RFC 4203 - Shared Risk Link Group (SRLG) sub-TLV
RFC 5185 OSPF Multi-Area Adjacency
RFC 3623 Graceful OSPF Restart — GR helper
RFC 3630 Traffic Engineering (TE) Extensions to OSPF Version 2
RFC 4203 for Shared Risk Link Group (SRLG) sub-TLV

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1772 Application of BGP in the Internet
RFC 1965 Confederations for BGP
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MD5
RFC 2439 BGP Route Flap Dampening
RFC 2547bis BGP/MPLS VPNs
RFC 2918 Route Refresh Capability for BGP-4
RFC 3107 Carrying Label Information in BGP-4
RFC 3392 Capabilities Advertisement with BGP4
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs)(previously RFC 2547bis BGP/MPLS VPNs)
RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP (previously RFC 1966 & 2796)
RFC 4724 Graceful Restart Mechanism for BGP – GR helper
RFC 4760 Multi-protocol Extensions for BGP
RFC 4893 BGP Support for Four-octet AS Number Space
RFC 5065 Confederations for BGP (obsoletes 3065)

IS-IS

RFC 1142 OSI IS-IS Intra-domain Routing Protocol (ISO 10589)
RFC 1195 Use of OSI IS-IS for routing in TCP/IP & dual environments
RFC 2763 Dynamic Hostname Exchange for IS-IS
RFC 2966 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 2973 IS-IS Mesh Groups
RFC 3373 Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies
RFC 3567 Intermediate System to Intermediate System (ISIS) Cryptographic Authentication
RFC 3719 Recommendations for Interoperable Networks using IS-IS
RFC 3784 Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)
RFC 3787 Recommendations for Interoperable IP Networks
RFC 3847 Restart Signaling for IS-IS – GR helper
RFC 4205 for Shared Risk Link Group (SRLG) TLV
draft-ietf-isis-igp-p2p-over-lan-05.txt

IPSec

RFC 2401 Security Architecture for the Internet Protocol
RFC 2409 The Internet Key Exchange (IKE)
RFC 3706 IKE Dead Peer Detection
RFC 3947 Negotiation of NAT-Traversal in the IKE
RFC 3948 UDP Encapsulation of IPsec ESP Packets
draft-ietf-ipsec-isakmp-xauth-06.txt – Extended Authentication within ISAKMP/Oakley (XAUTH)

draft-ietf-ipsec-isakmp-modecfg-05.txt – The ISAKMP Configuration Method

IPv6

RFC 1981 Path MTU Discovery for IPv6
RFC 2375 IPv6 Multicast Address Assignments
RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
RFC 2461 Neighbor Discovery for IPv6
RFC 2462 IPv6 Stateless Address Auto configuration
RFC 2463 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification
RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
RFC 2529 Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
RFC 2545 Use of BGP-4 Multiprotocol Extension for IPv6 Inter-Domain Routing
RFC 2710 Multicast Listener Discovery (MLD) for IPv6RFC 2740 OSPF for IPv6
RFC 3306 Unicast-Prefix-based IPv6 Multicast Addresses
RFC 3315 Dynamic Host Configuration Protocol for IPv6
RFC 3587 IPv6 Global Unicast Address Format
RFC3590 Source Address Selection for the Multicast Listener Discovery (MLD) Protocol
RFC 3810 Multicast Listener Discovery Version 2 (MLDv2) for IPv6
RFC 4007 IPv6 Scoped Address Architecture
RFC 4193 Unique Local IPv6 Unicast Addresses
RFC 4291 IPv6 Addressing Architecture
RFC 4552 Authentication/Confidentiality for OSPFv3
RFC 4659 BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
RFC 5072 IP Version 6 over PPP
RFC 5095 Deprecation of Type 0 Routing Headers in IPv6
draft-ietf-isis-ipv6-05
draft-ietf-isis-wg-multi-topology-xx.txt

Multicast

RFC 1112 Host Extensions for IP Multicasting (Snooping)
RFC 2236 Internet Group Management Protocol, (Snooping)
RFC 3376 Internet Group Management Protocol, Version 3 (Snooping)
RFC 2362 Protocol Independent Multicast-Sparse Mode (PIMSM)
RFC 3618 Multicast Source Discovery Protocol (MSDP)
RFC 3446 Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)
RFC 4601 Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)
RFC 4604 Using IGMPv3 and 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)
draft-ietf-pim-sm-bsr-06.txt
draft-rosen-vpn-mcast-15.txt Multicast in MPLS/BGP IP VPNs
draft-ietf-mboned-msdp-mib-01.txt
draft-ietf-l3vpn-2547bis-mcast-07: Multicast in MPLS/BGP IP VPNs
draft-ietf-l3vpn-2547bis-mcast-bgp-05: BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs
RFC 3956: Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address

MPLS — General

RFC 2430 A Provider Architecture DiffServ & TE
RFC 2474 Definition of the DS Field the IPv4 and IPv6 Headers (Rev)
RFC 2597 Assured Forwarding PHB Group (rev3260)
RFC 2598 An Expedited Forwarding PHB
RFC 3031 MPLS Architecture
RFC 3032 MPLS Label Stack Encoding
RFC 3443 Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks
RFC 4182 Removing a Restriction on the use of MPLS Explicit NULL
RFC 3140 Per-Hop Behavior Identification Codes
RFC 5332 MPLS Multicast Encapsulations

MPLS — LDP

RFC 3037 LDP Applicability
RFC 3478 Graceful Restart Mechanism for LDP – GR helper
RFC 5036 LDP Specification
RFC 5283 LDP extension for Inter-Area LSP
RFC 5443 LDP IGP Synchronization
draft-ietf-mpls-ldp-p2mp-05 LDP Extensions for Point-to-Multipoint and Multipoint-to-Multipoint LSP

MPLS/RSVP-TE

RFC 2702 Requirements for Traffic Engineering over MPLS
RFC2747 RSVP Cryptographic Authentication
RFC3097 RSVP Cryptographic Authentication
RFC 3209 Extensions to RSVP for Tunnels
RFC 3564 Requirements for Diff-Serv-aware TE
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 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 5151 Inter-domain MPLS and GMPLS Traffic Engineering – RSVP-TE Extensions

RFC 5712 MPLS Traffic Engineering Soft Preemption

draft-newton-mpls-te-dynamic-overbooking-00 A Diffserv-TE Implementation Model to dynamically change booking factors during failure events

RFC 5817 Graceful Shutdown in GMPLS Traffic Engineering Networks

MPLS — OAM

RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

draft-ietf-mpls-p2mp-lsp-ping-06 Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping

RIP

RFC 1058 RIP Version 1

RFC 2082 RIP-2 MD5 Authentication

RFC 2453 RIP Version 2

TCP/IP

RFC 768 UDP

RFC 1350 The TFTP Protocol (Rev.

RFC 791 IP

RFC 792 ICMP

RFC 793 TCP

RFC 826 ARP

RFC 854 Telnet

RFC 951 BootP (rev)

RFC 1519 CIDR

RFC 1542 Clarifications and Extensions for the Bootstrap Protocol

RFC 1812 Requirements for IPv4 Routers

RFC 2347 TFTP option Extension

RFC 2328 TFTP Blocksize Option

RFC 2349 TFTP Timeout Interval and Transfer

Size option

RFC 2401 Security Architecture for Internet Protocol

draft-ietf-bfd-mib-00.txtBidirectional Forwarding Detection Management Information Base

RFC 5880 Bidirectional Forwarding Detection

RFC 5881 BFD IPv4 and IPv6 (Single Hop)

RFC 5883 BFD for Multihop Paths

VRRP

RFC 2787 Definitions of Managed Objects for the Virtual Router Redundancy Protocol

RFC 3768 Virtual Router Redundancy Protocol

draft-ietf-vrrp-unified-spec-02: Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

PPP

RFC 1332 PPP IPCP

RFC 1377 PPP OSINLCP

RFC 1638/2878PPP BCP

RFC 1661 PPP (rev RFC2151)

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 2516 A Method for Transmitting PPP Over EthernetRFC 2615 PPP over SONET/SDH

RFC 2686 The Multi-Class Extension to Multi-Link PPP

Frame Relay

FRF.1.2 - PVC User-to-Network Interface (UNI) Implementation Agreement

FRF.5 - Frame Relay/ATM PVC Network Interworking Implementation

ANSI T1.617 Annex D, DSS1 — Signalling Specification For Frame Relay Bearer Service.

FRF2.2 -PVC Network-to- Network Interface (NNI) Implementation Agreement.

FRF.12 Frame Relay Fragmentation Implementation Agreement

FRF.16.1 Multilink Frame Relay UNI/ NNI Implementation Agreement

ITU-T Q.933 Annex A- Additional procedures for Permanent Virtual Connection (PVC) status management

ATM

RFC 1626 Default IP MTU for use over ATM AAL5

RFC 2514 Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management

RFC 2515 Definition of Managed Objects for ATM Management RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5

AF-TM-0121.000 Traffic Management Specification Version 4.1

ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/ 95

ITU-T Recommendation I.432.1 – BISDN user-network interface – Physical layer specification: General characteristics

GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3

GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1

AF-ILMI-0065.000 Integrated Local Management Interface (ILMI) Version 4.0

AF-TM-0150.00 Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR

AF-PHY-0086.001, Inverse Multiplexing for ATM (IMA) Specification Version 1.1

DHCP

RFC 2131 Dynamic HostConfiguration Protocol (REV)

RFC 3046 DHCP Relay Agent Information Option (Option 82)

RFC 1534 Interoperation between DHCP and BOOTP

VPLS

RFC 4762 Virtual Private LAN Services Using LDP

draft-ietf-l2vpn-vpls-mcast-reqts-04

draft-ietf-l2vpn-signaling-08

PSEUDOWIRE

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 3916 Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)

RFC 4717 Encapsulation Methods for Transport ATM over MPLS Networks (draft-ietf-pwe3-atm-encap-10.txt)

RFC 4816 PWE3 ATM Transparent Cell Transport Service (draft-ietf-pwe3-cell-transport-04.txt)

RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks (draft-ietf-pwe3-ethernet-encap-11.txt)

RFC 4619 Encapsulation Methods for Transport of Frame Relay over MPLS Networks (draft-ietf-pwe3-frame-relay-07.txt)

RFC 4446 IANA Allocations for PWE3

RFC 4447 Pseudowire Setup and Maintenance Using LDP (draft-ietf-pwe3-control-protocol-17.txt)

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

draft-ietf-l2vpn-vpws-iw-oam-02.txt, OAM Procedures for VPWS Interworking

draft-ietf-pwe3-oam-msg-map-14-txt, Pseudowire (PW) OAM Message Mapping

draft-ietf-l2vpn-arp-mediation-15.txt, ARP Mediation for IP Interworking of Layer 2 VPN

RFC6073, Segmented Pseudowire (draft-ietf-pwe3-segmented-pw-18.txt)

draft-ietf-pwe3-dynamic-ms-pw-13.txt , Dynamic Placement of Multi Segment Pseudo Wires

draft-ietf-pwe3-redundancy-bit-03.txt, Pseudowire Preferential Forwarding Status bit definition

draft-ietf-pwe3-redundancy-03.txt, Pseudowire (PW) Redundancy

draft-ietf-pwe3-fat-pw-05 Flow Aware Transport of Pseudowires over an MPLS PSN

MFA Forum 9.0.0 The Use of Virtual trunks for ATM/MPLS Control Plane 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

ANCP/L2CP

RFC5851 ANCP framework

draft-ietf-ancp-protocol-02.txt ANCP Protocol

Voice /Video Performance

ITU-T G.107 The E Model- A computational model for use in planning.

ETSI TS 101 329-5 Annex E extensions- QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring

ITU-T Rec. P.564 - Conformance testing for voice over IP transmission quality assessment models

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.

RFC 3550 Appendix A.8- RTP: A Transport Protocol for Real-Time Applications- Estimating the Interarrival Jitter

CIRCUIT EMULATION

RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)

MEF-8 Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks, October 2004

RFC 5287 Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks

SONET/SDH

ITU-G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

RADIUS

RFC 2865 Remote Authentication Dial In User Service

RFC 2866 RADIUS Accounting

SSH

draft-ietf-secsh-architecture.txtSSH Protocol Architecture

draft-ietf-secsh-userauth.txt SSH Authentication Protocol

draft-ietf-secsh-transport.txt SSH Transport Layer Protocol

draft-ietf-secsh-connection.txt SSH Connection Protocol

draft-ietf-secsh- newmodes.txt SSH Transport Layer Encryption Modes

TACACS+

draft-grant-tacacs-02.txt

Timing

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

ITU-T G.781 Telecommunication Standardization Section of ITU, Synchronization layer functions, issued 09/2008

ITU-T G.813 Telecommunication Standardization Section of ITU, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003.

GR-1244-CORE Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005

ITU-T G.8261 Telecommunication Standardization Section of ITU, Timing and synchronization aspects in packet networks, issued 04/2008.

ITU-T G.8262 Telecommunication Standardization Section of ITU, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007.

ITU-T G.8264 Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008.

NETWORK MANAGEMENT

ITU-T X.721: Information technology-OSI-Structure of Management Information

ITU-T X.734: Information technology-OSI-Systems Management: Event Report Management Function

M.3100/3120 Equipment and Connection Models

TMF 509/613 Network Connectivity Model

RFC 1157 SNMPv1

RFC 1215 A Convention for Defining Traps for use with the SNMP

RFC 1657 BGP4-MIB

RFC 1724 RIPv2-MIB

RFC 1850 OSPF-MIB

RFC 1907 SNMPv2-MIB

RFC 2011 IP-MIB

RFC 2012 TCP-MIB

RFC 2013 UDP-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2452 IPv6 Management Information Base for the Transmission Control Protocol

RFC 2454 IPv6 Management Information Base for the User Datagram Protocol

RFC 2465 Management Information Base for IPv6: Textual Conventions and General Group

RFC 2558 SONET-MIB

RFC 2571 SNMP-Framework MIB

RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-Target-&-notification-MIB

RFC 2574 SNMP-User-based-SMMIB

RFC 2575 SNMP-View-based ACM-MIB

RFC 2576 SNMP-Community-MIB

RFC 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 Inverted-stack-MIB

RFC 2987 VRRP-MIB

RFC 3014 Notification-log MIB

RFC 3019 IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol

RFC 3164 Syslog

RFC 3273 HCRMON-MIB

RFC 3411 An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks

RFC 3412 Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)

RFC 3413 Simple Network Management Protocol (SNMP) Applications

RFC 3414 User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)

RFC 3418 SNMP MIB

RFC 4292 IP-Forward-MIB

RFC 4293 MIB for the Internet Protocol

RFC 5101 Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information

draft-ietf-ospf-mib-update-04.txt

draft-ietf-mpls-lsr-mib-06.txt

draft-ietf-mpls-te-mib-04.txt

draft-ietf-mpls-ldp-mib-07.txt

draft-ietf-isis-wg-mib-05.txt

IANA-IFType-MIB

IEEE8023-LAG-MIB

Proprietary MIBs

TIMETRA-APS-MIB.mib

TIMETRA-ATM-MIB.mib

TIMETRA-BGP-MIB.mib

TIMETRA-BSX-NG-MIB.mib

TIMETRA-CAPABILITY-7750-V4v0.mib

TIMETRA-CFLOWD-MIB.mib

TIMETRA-CHASSIS-MIB.mib

TIMETRA-CLEAR-MIB.mib

TIMETRA-FILTER-MIB.mib

TIMETRA-GLOBAL-MIB.mib

TIMETRA-IGMP-MIB.mib

TIMETRA-ISIS-MIB.mib

TIMETRA-LAG-MIB.mib

TIMETRA-LDP-MIB.mib

TIMETRA-LOG-MIB.mib

TIMETRA-MIRROR-MIB.mib

TIMETRA-MPLS-MIB.mib

TIMETRA-NG-BGP-MIB.mib

TIMETRA-OAM-TEST-MIB.mib

TIMETRA-OSPF-NG-MIB.mib

TIMETRA-OSPF-V3-MIB.mib

TIMETRA-PIM-NG-MIB.mib

TIMETRA-PORT-MIB.mib

TIMETRA-PPP-MIB.mib

TIMETRA-QOS-MIB.mib

TIMETRA-RIP-MIB.mib

TIMETRA-ROUTE-POLICY-MIB.mib

TIMETRA-RSVP-MIB.mib

TIMETRA-SECURITY-MIB.mib

TIMETRA-SERV-MIB.mib

TIMETRA-SUBSCRIBER-MGMTMIB.mib

TIMETRA-SYSTEM-MIB.mib

TIMETRA-TC-MIB.mib

TIMETRA-VRRP-MIB.mib

TIMETRA-VRTR-MIB.mib

Index