Alcatel-Lucent 7705

SERVICE AGGREGATION ROUTER OS | RELEASE 4.0

OAM AND DIAGNOSTICS GUIDE

Alcatel-Lucent assumes no responsibility for the accuracy of the information presented, which is subject to change without notice.

Disclaimers

Alcatel-Lucent products are intended for commercial uses. Without the appropriate network design engineering, they must not be sold, licensed or otherwise distributed for use in any hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life-support machines, or weapons systems, in which the failure of products could lead directly to death, personal injury, or severe physical or environmental damage. The customer hereby agrees that the use, sale, license or other distribution of the products for any such application without the prior written consent of Alcatel-Lucent, shall be at the customer's sole risk. The customer hereby agrees to defend and hold Alcatel-Lucent harmless from any claims for loss, cost, damage, expense or liability that may arise out of or in connection with the use, sale, license or other distribution of the products in such applications.

This document may contain information regarding the use and installation of non-Alcatel-Lucent products. Please note that this information is provided as a courtesy to assist you. While Alcatel-Lucent tries to ensure that this information accurately reflects information provided by the supplier, please refer to the materials provided with any non-Alcatel-Lucent product and contact the supplier for confirmation. Alcatel-Lucent assumes no responsibility or liability for incorrect or incomplete information provided about non-Alcatel-Lucent products.

However, this does not constitute a representation or warranty. The warranties provided for Alcatel-Lucent products, if any, are set forth in contractual documentation entered into by Alcatel-Lucent and its customers.

This document was originally written in English. If there is any conflict or inconsistency between the English version and any other version of a document, the English version shall prevail.

Table of Contents

List of Tables

List of Figures

List of Acronyms

Acronym	Expansion
2G	second generation wireless telephone technology
3DES	triple DES (data encryption standard)
3G	third generation mobile telephone technology
5620 SAM	5620 Service Aware Manager
7705 SAR	7705 Service Aggregation Router
7710 SR	7710 Service Router
7750 SR	7750 Service Router
9500 MPR	9500 Microwave Packet Radio
ABR	available bit rate area border router
AC	alternating current attachment circuit
ACK	acknowledge
ACL	access control list
ACR	adaptive clock recovery
ADP	automatic discovery protocol
AFI	authority and format identifier
AIS	alarm indication signal
ANSI	American National Standards Institute
Apipe	ATM VLL
APS	automatic protection switching
ARP	address resolution protocol
A/S	active/standby
AS	autonomous system

Acronym	Expansion
ASAP	any service, any port
ASBR	autonomous system boundary router
ASN	autonomous system number
ATM	asynchronous transfer mode
ATM PVC	ATM permanent virtual circuit
B3ZS	bipolar with three-zero substitution
Batt A	battery A
B-bit	beginning bit (first packet of a fragment)
Bellcore	Bell Communications Research
BFD	bidirectional forwarding detection
BGP	border gateway protocol
BITS	building integrated timing supply
BMCA	best master clock algorithm
BMU	broadcast, multicast, and unknown traffic Traffic that is not unicast. Any nature of multipoint traffic: • broadcast (that is, all 1s as the destination IP to represent all destinations within the subnet) • multicast (that is, traffic typically identified by the destination address, uses special destination address); for IP, the destination must be 224.0.0.0 to 239.255.255.255 • unknown (that is, the destination is typically a valid unicast address but the destination port/interface is not yet known; therefore, traffic needs to be forwarded to all destinations; unknown traffic is treated as broadcast)
BOF	boot options file
BPDU	bridge protocol data unit
BRAS	Broadband Remote Access Server
BSC	Base Station Controller
BSTA	Broadband Service Termination Architecture

Acronym	Expansion
BTS	base transceiver station
CAS	channel associated signaling
CBN	common bonding networks
CBS	committed buffer space
CC	control channel continuity check
CCM	continuity check message
CE	customer edge circuit emulation
CEM	circuit emulation
CES	circuit emulation services
CESoPSN	circuit emulation services over packet switched network
CFM	connectivity fault management
CIDR	classless inter-domain routing
CIR	committed information rate
CLI	command line interface
CLP	cell loss priority
CoS	class of service
CPE	customer premises equipment
Cpipe	circuit emulation (or TDM) VLL
CPM	Control and Processing Module (CPM is used instead of CSM when referring to CSM filtering to align with CLI syntax used with other SR products). CSM management ports are referred to as CPM management ports in the CLI.
CPU	central processing unit
CRC	cyclic redundancy check
CRON	a time-based scheduling service (from chronos = time)

Acronym	Expansion
CSM	Control and Switching Module
CSNP	complete sequence number PDU
CSPF	constrained shortest path first
C-TAG	customer VLAN tag
CV	connection verification customer VLAN (tag)
CW	control word
DC	direct current
DC-C	DC return - common
DCE	data communications equipment
DC-I	DC return - isolated
DCO	digitally controlled oscillator
DDoS	distributed DoS
DES	data encryption standard
DF	do not fragment
DHB	decimal, hexadecimal, or binary
DHCP	dynamic host configuration protocol
DHCPv6	dynamic host configuration protocol for IPv6
DIS	designated intermediate system
DM	delay measurement
DNS	domain name server
DoS	denial of service
dot1p	IEEE 802.1p bits, found in Ethernet or VLAN ingress packet headers and used to map traffic to up to eight forwarding classes
dot1q	IEEE 802.1q encapsulation for Ethernet interfaces
DPI	deep packet inspection

Acronym	Expansion
DPLL	digital phase locked loop
DSCP	differentiated services code point
DSL	digital subscriber line
DSLAM	digital subscriber line access multiplexer
DTE	data termination equipment
DU	downstream unsolicited
DUID	DHCP unique identifier
DV	delay variation
e911	enhanced 911 service
EAP	Extensible Authentication Protocol
EAPOL	EAP over LAN
E-bit	ending bit (last packet of a fragment)
ECMP	equal cost multi-path
EFM	Ethernet in the first mile
EGP	exterior gateway protocol
EIA/TIA-232	Electronic Industries Alliance/Telecommunications Industry Association Standard 232 (also known as RS-232)
ELER	egress label edge router
E&M	ear and mouth earth and magneto exchange and multiplexer
Epipe	Ethernet VLL
EPL	Ethernet private line
ERO	explicit route object
ESD	electrostatic discharge
ESMC	Ethernet synchronization message channel
ETE	end-to-end

Acronym	Expansion
ETH-CFM	Ethernet connectivity fault management (IEEE 802.1ag)
EVDO	evolution - data optimized
EVPL	Ethernet virtual private link
EXP bits	experimental bits (currently known as TC)
FC	forwarding class
FCS	frame check sequence
FDB	forwarding database
FDL	facilities data link
FEAC	far-end alarm and control
FEC	forwarding equivalence class
FF	fixed filter
FIB	forwarding information base
FIFO	first in, first out
FNG	fault notification generator
FOM	figure of merit
FRR	fast reroute
FTN	FEC-to-NHLFE
FTP	file transfer protocol
GFP	generic framing procedure
GigE	Gigabit Ethernet
GRE	generic routing encapsulation
GSM	Global System for Mobile Communications (2G)
HCM	high capacity multiplexing
HDB3	high density bipolar of order 3
HEC	header error control
HMAC	hash message authentication code

Acronym	Expansion
HSDPA	high-speed downlink packet access
HSPA	high-speed packet access
HVPLS	hierarchical virtual private line service
IANA	internet assigned numbers authority
IBN	isolated bonding networks
ICMP	Internet control message protocol
ICMPv6	Internet control message protocol for IPv6
ICP	IMA control protocol cells
IEEE	Institute of Electrical and Electronics Engineers
IEEE 1588v2	Institute of Electrical and Electronics Engineers standard 1588-2008
IES	Internet Enhanced Service
IETF	Internet Engineering Task Force
IGP	interior gateway protocol
ILER	ingress label edge router
ILM	incoming label map
IMA	inverse multiplexing over ATM
IOM	input/output module
IP	Internet Protocol
IPCP	Internet Protocol Control Protocol
IPIP	IP in IP
Ipipe	IP interworking VLL
IPoATM	IP over ATM
IS-IS	Intermediate System-to-Intermediate System
IS-IS-TE	IS-IS-traffic engineering (extensions)
ISO	International Organization for Standardization

Acronym	Expansion
LB	loopback
lbf-in	pound force inch
LBM	loopback message
LBO	line buildout
LBR	loopback reply
LCP	link control protocol
LDP	label distribution protocol
LER	label edge router
LFIB	label forwarding information base
LIB	label information base
LLDP	link layer discovery protocol
LLDPDU	link layer discovery protocol data unit
LLF	link loss forwarding
LLID	loopback location ID
LM	loss measurement
LSA	link-state advertisement
LSDB	link-state database
LSP	label switched path link-state PDU (for IS-IS)
LSR	label switch router link-state request
LSU	link-state update
LT	linktrace
LTE	line termination equipment
LTM	linktrace message
LTN	LSP ID to NHLFE

Acronym	Expansion
LTR	linktrace reply
MA	maintenance association
MAC	media access control
MA-ID	maintenance association identifier
MBB	make-before-break
MBS	maximum buffer space maximum burst size media buffer space
MBSP	mobile backhaul service provider
MC-MLPPP	multi-class multilink point-to-point protocol
MD	maintenance domain
MD5	message digest version 5 (algorithm)
MDA	media dependent adapter
MDDB	multidrop data bridge
MDL	maintenance data link
ME	maintenance entity
MED	multi-exit discriminator
MEF	Metro Ethernet Forum
MEG	maintenance entity group
MEG-ID	maintenance entity group identifier
MEN	Metro Ethernet network
MEP	maintenance association end point
MFC	multi-field classification
MHF	MIP half function
MIB	management information base
MIP	maintenance association intermediate point

Acronym	Expansion
MIR	minimum information rate
MLPPP	multilink point-to-point protocol
MP	merge point multilink protocol
MP-BGP	multiprotocol border gateway protocol
MPLS	multiprotocol label switching
MPR	see 9500 MPR
MRRU	maximum received reconstructed unit
MRU	maximum receive unit
MSDU	MAC Service Data Unit
MS-PW	multi-segment pseudowire
MTIE	maximum time interval error
MTSO	mobile trunk switching office
MTU	maximum transmission unit multi-tenant unit
M-VPLS	management virtual private line service
MW	microwave
N·m	newton meter
NBMA	non-broadcast multiple access (network)
NE	network element
NET	network entity title
NHLFE	next hop label forwarding entry
NHOP	next-hop
NLRI	network layer reachability information
NNHOP	next next-hop
NNI	network-to-network interface

Acronym	Expansion
Node B	similar to BTS but used in 3G networks — term is used in UMTS (3G systems) while BTS is used in GSM (2G systems)
NSAP	network service access point
NSSA	not-so-stubby area
NTP	network time protocol
OAM	operations, administration, and maintenance
OAMPDU	OAM protocol data units
OC3	optical carrier, level 3
ORF	outbound route filtering
OS	operating system
OSI	Open Systems Interconnection (reference model)
OSINLCP	OSI Network Layer Control Protocol
OSPF	Open Shortest Path First
OSPF-TE	OSPF-traffic engineering (extensions)
OSS	operations support system
OSSP	Organization Specific Slow Protocol
OTP	one time password
PADI	PPPoE active discovery initiation
PADR	PPPoE active discovery request
PAE	port authentication entities
PCP	priority point code
PDU	protocol data units
PDV	packet delay variation
PDVT	packet delay variation tolerance
PE	provider edge router
PHB	per-hop behavior

Acronym	Expansion
PHY	physical layer
PID	protocol ID
PIR	peak information rate
PLCP	Physical Layer Convergence Protocol
PLR	point of local repair
POP	point of presence
POS	packet over SONET
PPP	point-to-point protocol
PPPoE	point-to-point protocol over Ethernet
PRC	primary reference clock
PSN	packet switched network
PSNP	partial sequence number PDU
PTP	precision time protocol performance transparency protocol
PVC	permanent virtual circuit
PVCC	permanent virtual channel connection
PW	pseudowire
PWE	pseudowire emulation
PWE3	pseudowire emulation edge-to-edge
QL	quality level
QoS	quality of service
RADIUS	Remote Authentication Dial In User Service
RAN	Radio Access Network
RBS	robbed bit signaling
RD	route distinguisher
RDI	remote defect indication

Acronym	Expansion
RED	random early discard
RESV	reservation
RIB	routing information base
RJ-45	registered jack 45
RNC	Radio Network Controller
RRO	record route object
RS-232	Recommended Standard 232 (also known as EIA/TIA-232)
RSHG	residential split horizon group
RSTP	Rapid Spanning Tree Protocol
RSVP-TE	resource reservation protocol - traffic engineering
RT	receive/transmit
RTM	routing table manager
RTN	battery return
RTP	real-time protocol
R&TTE	Radio and Telecommunications Terminal Equipment
RTU	remote terminal unit
RU	rack unit
SAA	service assurance agent
SAP	service access point
SAR-8	7705 Service Aggregation Router - 8-slot chassis
SAR-18	7705 Service Aggregation Router - 18-slot chassis
SAR-F	7705 Service Aggregation Router - fixed form-factor chassis
SAToP	structure-agnostic TDM over packet
SCADA	surveillance, control and data acquisition
SCP	secure copy
SD	signal degrade

Acronym	Expansion
SDH	synchronous digital hierarchy
SDI	serial data interface
SDP	service destination point
SE	shared explicit
SF	signal fail
SFP	small form-factor pluggable (transceiver)
SGT	self-generated traffic
SHA-1	secure hash algorithm
SHG	split horizon group
SIR	sustained information rate
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SNPA	subnetwork point of attachment
SNTP	simple network time protocol
SONET	synchronous optical networking
S-PE	switching provider edge router
SPF	shortest path first
SPT	shortest path tree
SR	service router (includes 7710 SR, 7750 SR)
SRLG	shared risk link group
SSH	secure shell
SSM	synchronization status messaging
SSU	system synchronization unit
S-TAG	service VLAN tag
STM1	synchronous transport module, level 1
SVC	switched virtual circuit

Acronym	Expansion
SYN	synchronize
TACACS+	Terminal Access Controller Access-Control System Plus
TC	traffic class (formerly known as EXP bits)
TCP	transmission control protocol
TDEV	time deviation
TDM	time division multiplexing
TE	traffic engineering
TFTP	trivial file transfer protocol
TLDP	targeted LDP
TLV	type length value
ToS	type of service
T-PE	terminating provider edge router
TPID	tag protocol identifier
TPMR	two-port MAC relay
TTL	time to live
TTM	tunnel table manager
U-APS	unidirectional automatic protection switching
UBR	unspecified bit rate
UDP	user datagram protocol
UMTS	Universal Mobile Telecommunications System (3G)
UNI	user-to-network interface
V.35	V-series Recommendation 35
VC	virtual circuit
VCC	virtual channel connection
VCCV	virtual circuit connectivity verification
VCI	virtual circuit identifier

Acronym	Expansion
VID	VLAN ID
VLAN	virtual LAN
VLL	virtual leased line
VoIP	voice over IP
Vp	peak voltage
VP	virtual path
VPC	virtual path connection
VPI	virtual path identifier
VPLS	virtual private LAN service
VPN	virtual private network
VPRN	virtual private routed network
VRF	virtual routing and forwarding table
VSE	vendor-specific extension
VSO	vendor-specific option
WCDMA	wideband code division multiple access (transmission protocol used in UMTS networks)
WRED	weighted random early discard
WTR	wait to restore

Preface

About This Guide

This guide describes Operations, Administration and Management (OAM) and diagnostic tools provided by the 7705 SAR 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 guide is intended for network administrators who are responsible for configuring the 7705 SAR routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Protocols, standards, and services described in this guide include the following:

- CLI concepts
- operations, administration, and maintenance (OAM) operations

List of Technical Publications

The 7705 SAR OS documentation set is composed of the following guides:

- 7705 SAR OS Basic System Configuration Guide

  This guide describes basic system configurations and operations.
- 7705 SAR OS System Management Guide

  This guide describes system security and access configurations as well as event logging and accounting logs.
- 7705 SAR OS Interface Configuration Guide

  This guide describes card and port provisioning.

- 7705 SAR OS Router Configuration Guide

  This guide describes logical IP routing interfaces, IP-based filtering, and routing policies.
- 7705 SAR OS MPLS Guide

  This guide describes how to configure Multiprotocol Label Switching (MPLS), Resource Reservation Protocol for Traffic Engineering (RSVP-TE), and Label Distribution Protocol (LDP).
- 7705 SAR OS Services Guide

  This guide describes how to configure service parameters such as service access points (SAPs), service destination points (SDPs), customer information, and user services.
- 7705 SAR OS Quality of Service Guide

  This guide describes how to configure Quality of Service (QoS) policy management.
- 7705 SAR OS Routing Protocols Guide

  This guide provides an overview of dynamic routing concepts and describes how to configure them.
- 7705 SAR OS OAM and Diagnostics Guide

  This guide provides information on Operations, Administration and Maintenance (OAM) tools.

Multiple PDF File Search

You can use Adobe Reader, Release 6.0 or later, to search multiple PDF files for a term. Adobe Reader displays the results in a display panel. The results are grouped by PDF file. You can expand the entry for each file.

Note: The PDF files in which you search must be in the same folder.

To search multiple PDF files for a term:

Step 1. Open Adobe Reader.

Step 2. Choose Edit – Search from the Adobe Reader main menu. The Search panel appears.

Step 3. Enter the term to search for.

Step 4. Select the All PDF Documents in radio button.

Step 5. Choose the folder in which to search using the drop-down menu.

Step 6. Select the following criteria if required:

- Whole words only
- Case-Sensitive
- Include Bookmarks
- Include Comments

Step 7. Click on the Search button.

Adobe Reader displays the search results. You can expand the entries for each file by clicking on the + symbol.

Step 8. Click on a search result to go directly to that location in the selected file.

Technical Support

If you purchased a service agreement for your 7705 SAR 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, check this link for instructions to contact Support personnel:

Web: http://support.alcatel-lucent.com

Getting Started

In This Chapter

This chapter provides the process flow information required to configure Operations, Administration and Management (OAM) tools.

Alcatel-Lucent 7705 SAR OAM Configuration Process

Table 1 lists the tasks necessary to 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	Reference
Diagnostics/Service verification	OAM	OAM and SAA on page 33 Tools on page 197
Reference	List of IEEE, IETF, and other proprietary entities	Standards and Protocol Support on page 229

Notes on 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F

The 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F run the same operating system software. The main difference between the products is their hardware platforms.

The 7705 SAR-8 is an 8-slot chassis that supports 2 CSMs, a Fan module, and 6 adapter cards. The 7705 SAR-18 chassis has 18 slots; in Release 4.0, it supports 2 CSMs, a Fan module, an Alarm module, and 12 adapter cards.

The 7705 SAR-F chassis has a fixed hardware configuration. The 7705 SAR-F replaces the CSM, Fan module, and the 16-port T1/E1 ASAP Adapter card and 8-port Ethernet Adapter card with an all-in-one unit that provides comparable functional blocks, as detailed in Table 2.

The fixed configuration of the 7705 SAR-F means that provisioning the router at the "card slot" and "type" levels is preset and is not user-configurable. Operators begin configurations at the port level.

Note: Unless stated otherwise, references to the terms "Adapter card" and "CSM" throughout the 7705 SAR OS documentation set include the equivalent functional blocks on the 7705 SAR-F.

Table 2: 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F Comparison

7705 SAR-8, 7705 SAR-18	7705 SAR-F	Notes
CSM	Control and switching functions	The control and switching functions include the console and management interfaces, the alarm and fan functions, the synchronization interfaces, system LEDs, and so on.
Fan module	Integrated with the control and switching functions	

Table 2: 7705 SAR-8, 7705 SAR-18, and 7705 SAR-F Comparison (Continued)

7705 SAR-8, 7705 SAR-18	7705 SAR-F	Notes
16-port T1/E1 ASAP Adapter card	16 individual T1/E1 ports on the faceplate	The T1/E1 ports on the 7705 SAR-F are equivalent to the T1/E1 ports on the 16-port T1/E1 ASAP Adapter card, version 1, except that the 16 T1/E1 ports on the 7705 SAR-F support multiple synchronization sources to support two timing references. The 16-port T1/E1 ASAP Adapter card, version 2, also supports two timing references. On the 7705 SAR-8 and 7705 SAR-18, the CLI indicates the MDA type for the 16-port T1/E1 ASAP Adapter card as `a16-chds1` for version 1 `and a16-chds1v2` for version 2. On the 7705 SAR-F, the CLI indicates the MDA type for the 7705 SAR-F ports as `i16-chds1.`
8-port Ethernet Adapter card	8 individual Ethernet ports on the faceplate	The –48 VDC versions of the 7705 SAR-8 support two versions of the 8-port Ethernet Adapter card, with version 2 having additional support for Synchronous Ethernet. The +24 VDC version of the 7705 SAR-8 supports only version 2 of the 8-port Ethernet Adapter card. The 7705 SAR-18 supports only version 2 of the card. The Ethernet ports on the 7705 SAR-F are functionally equivalent to the Ethernet ports on version 2 of the 8-port Ethernet Adapter card and support multiple synchronization sources to support two timing references. On the 7705 SAR-8, the CLI indicates the MDA type for the 8-port Ethernet Adapter card as `a8-eth` or `a8-ethv2`. On the 7705 SAR-18, the CLI indicates the MDA type as `a8-ethv2`. On the 7705 SAR-F, the CLI indicates the MDA type for the 7705 SAR-F Ethernet ports as `i8-eth`.
Requires user configuration at card (IOM) and MDA (adapter card) levels	Configuration at card (IOM) and MDA (adapter card) levels is preset and users cannot change these types	

OAM and SAA

In This Chapter

This chapter provides information about the Operations, Administration and Maintenance (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 that 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 OAM tools is provided, with the ability to test each of the individual packet operations.

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

The suite of OAM diagnostics supplements the basic IP ping and traceroute operations with diagnostics specialized for the different levels in the service delivery model. In addition, there are diagnostics for MPLS LSPs, SDPs, and Services within a service.

ICMP and ICMPv6 Diagnostics

Internet Control Message Protocol (ICMP) is part of the IP suite as defined in RFC 792, *Internet Control Message Protocol*, for IPv4 and RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*.

ICMP and ICMPv6 send and receive control and error messages used to manage the behavior of the TCP/IP stack. ICMP and ICMPv6 provide:

- debugging tools and error reporting mechanisms to assist in troubleshooting an IP network
- the ability to send and receive error and control messages to far-end IP entities

Ping

Ping is used to determine if there is IP layer connectivity between the 7705 SAR and another node in the network.

Traceroute

Traceroute is used to determine the path that an IP packet takes from the 7705 SAR to a specified router.

LSP Diagnostics

The 7705 SAR LSP diagnostics are implementations of LSP ping and LSP traceroute based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. 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.

LSP Ping

LSP ping, as described in RFC 4379, provides a mechanism to detect data plane failures in MPLS LSPs. For a given FEC, LSP ping verifies whether the packet reaches the egress label edge router (ELER).

LSP Traceroute

LSP traceroute sends a packet to each transit LSR along a communications path until the far-end router is reached. The path is traced one LSR at a time, where each LSR that receives a traceroute packet replies to the initiating 7705 SAR with a packet that identifies itself. Once the final LSR is identified, the initiating LSR has a list of all LSRs on the path. Like IP traceroute, LSP traceroute is a hop-by-hop operation (that is, LSR by LSR).

Use LSP traceroute to determine the exact litigation of LSP failures.

SDP Diagnostics

The 7705 SAR SDP diagnostics include SDP ping and SDP MTU path discovery.

SDP Ping

SDP ping performs in-band unidirectional 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 unidirectional test, SDP ping tests:

- the egress SDP ID encapsulation
- the ability to reach the far-end IP address of the SDP ID within the SDP encapsulation
- the path MTU to the far-end IP address over the SDP ID
- the 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 unidirectional tunnels, the remote SDP ID must be specified and must exist as a configured SDP ID on the far-end 7705 SAR. SDP round-trip testing is an extension of SDP connectivity testing with the additional ability to test:

- the remote SDP ID encapsulation
- the potential service round-trip time
- the round-trip path MTU
- the 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 their interfaces. The largest packet (including headers) can be as large as the Maximum Transmission Unit (MTU). An MTU specifies the largest packet size, measured in octets, that can be transmitted through a network entity. It is important to understand the MTU of the entire path (end-to-end) when provisioning services, especially for VLL services where the service must support the ability to transmit the extra large customer packets.

The Path MTU Discovery tool is a powerful tool that enables service providers to get the exact MTU supported between the service ingress and service termination points, accurate to 1 byte.

Service Diagnostics

The Alcatel-Lucent 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

Service (SVC) ping is initiated from a 7705 SAR router to verify round-trip connectivity and delay to the far end of the service. The Alcatel-Lucent implementation of Service ping applies to GRE, IP, 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

Note: By default, service ping uses GRE encapsulation.

VLL Diagnostics

This section describes VCCV (Virtual Circuit Connectivity Verification) ping and VCCV trace, the VLL diagnostic capabilities for the 7705 SAR.

VCCV Ping

VCCV ping is used to check the connectivity (in-band) of a VLL. It checks that the destination (target) PE is the egress point 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, GRE, or IP SDP.

VCCV Ping Application

VCCV creates an IP control channel within the pseudowire between PE1 and PE2 (see Figure 1). PE2 should be able to distinguish, on the receive side, VCCV control messages from user packets on that VLL.

Figure 1: VCCV Ping Application

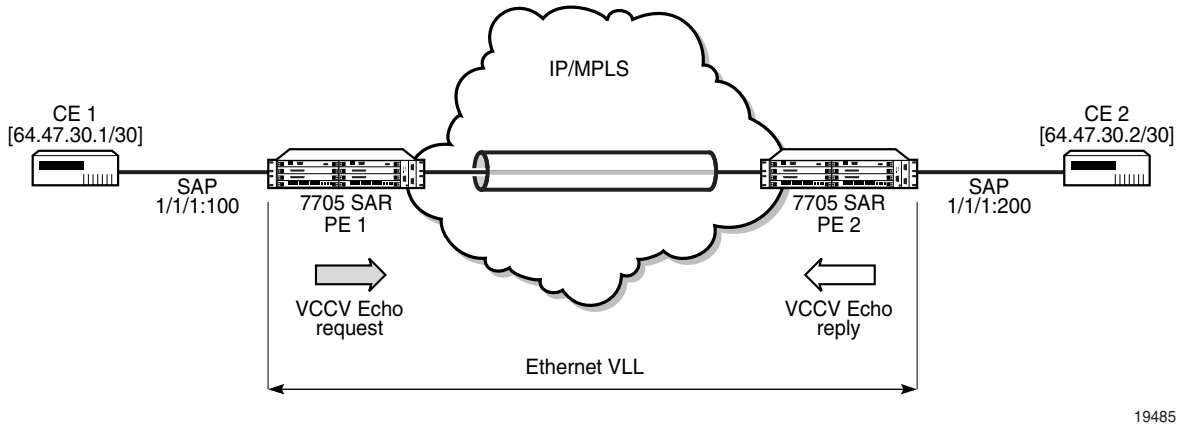

VCCV-based pseudowire (PW) tests are only supported on dynamically signaled PWs (not on statically signaled PWs).

There are three methods of encapsulating a VCCV message in a VLL, which translates into three types of control channels, as follows:

- Type 1 — in-band VCCV (special control word)

  Type 1 uses the OAM control word, which is shown in Figure 2.

Figure 2: OAM Control Word Format

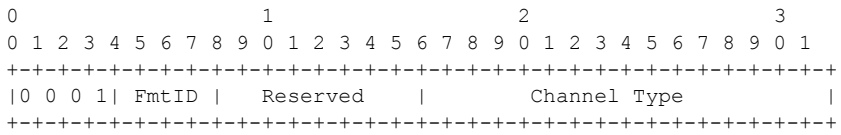

  In Figure 2, the first nibble is set to 0x1. The Format ID and the Reserved fields are set to 0 and the Channel Type is the code point associated with the VCCV IP control channel, as specified in the PWE3 IANA registry [RFC 4446]. The channel type value of 0x21 indicates that the Associated Channel carries an IPv4 packet.

  The use of the OAM control word assumes that the *draft-martini* control word is also used for the user packets. This means that if the control word is optional for a VLL and is not configured, the 7705 SAR PE node will only advertise the router alert label as the CC capability in the Label Mapping message.

  This method is supported by the 7705 SAR.

- Type 2 — out-of-band VCCV (router alert above the service label)

  The 7705 SAR uses the router alert label immediately above the VC label to identify the VCCV ping message. This method has a drawback in that if ECMP is applied to the outer LSP label, such as the 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 7705 SAR when a 7750 SR node acts as an LSR in the core of the network. If a 7705 SAR acts as an LSR in the core of the network, the VCCV type 2 message will instead follow the data path.

- Type 3 — TTL expiry VCCV (service label TTL = 1 and special control word)

  This method is not supported by the 7705 SAR.

When sending the label mapping message for the VLL, PE1 and PE2 must indicate which of the above OAM packet encapsulation methods (that is, which control channel type) they support. This is accomplished by including an optional VCCV TLV in the PW FEC interface parameter field. The format of the VCCV TLV is shown in Figure 3.

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

Figure 3: VCCV TLV

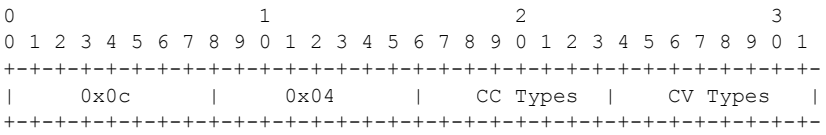

In Figure 3, the Control Channel (CC) Type field is a bit mask used to indicate if the PE supports none, one, or many control channel types:

- 0x00 — none of the following VCCV control channel types (Type 1, Type 2, or Type 3) are supported
- 0x01 — (Type 1, in-band) PWE3 OAM control word (see Figure 2)
- 0x02 — (Type 2, out-of-band) MPLS router alert label
- 0x04 — (Type 3, not supported on the 7705 SAR) MPLS inner label TTL = 1

If both PE nodes support more than one of the CC types, then a 7705 SAR PE will make use of the CC type with the lowest type value. For instance, OAM control word (0x01) will be used in preference to the MPLS router alert label (0x02).

The Connectivity Verification (CV) Type field is a bit mask used to indicate the specific type of VCCV packets to be sent over the VCCV control channel. The possible values supported on the 7705 SAR are:

- 0x00 — none of the following VCCV packet types are supported
- 0x02 — LSP ping

  This value (0x02) is used in the VCCV ping application and applies to a VLL over an MPLS, GRE, or IP SDP.

A VCCV ping is an LSP echo request message as defined in RFC 4379. It contains a Layer 2 FEC stack TLV in which it must include the sub-TLV type 10 FEC 128 pseudowire. It also contains a field that indicates to the destination PE which reply mode to use:

- do not reply

  This mode is supported by the 7705 SAR.
- reply by an IPv4 UDP packet

  This mode is supported by the 7705 SAR.
- reply via an IPv4 UDP packet with router alert

  This mode is not supported by the 7705 SAR.

  Note: Do not confuse this mode, which sets the router alert bit in the IP header, with the CC type that makes use of the router alert label.

- reply by application-level control channel

  This mode sends the reply message in-band 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 7705 SAR.

The VCCV ping reply has the same format as an LSP echo reply message as defined in RFC 4379. The message is sent via the reply mode requested by PE1. The return codes supported are the same as those currently supported in the 7705 SAR LSP ping capability.

The VCCV ping feature is in addition to the service ping OAM feature that can be used to test a service between 7705 SAR nodes. The VCCV ping feature can test connectivity of a VLL with any third-party node that is compliant with RFC 5085.

VCCV Ping in a Multi-Segment Pseudowire

Figure 4 displays an example of an application of VCCV ping over a multi-segment pseudowire (MS-PW). Pseudowire switching provides the user with the ability to create a VLL service by cross-connecting two spoke SDPs.

Figure 4: VCCV Ping Over a Multi-Segment Pseudowire

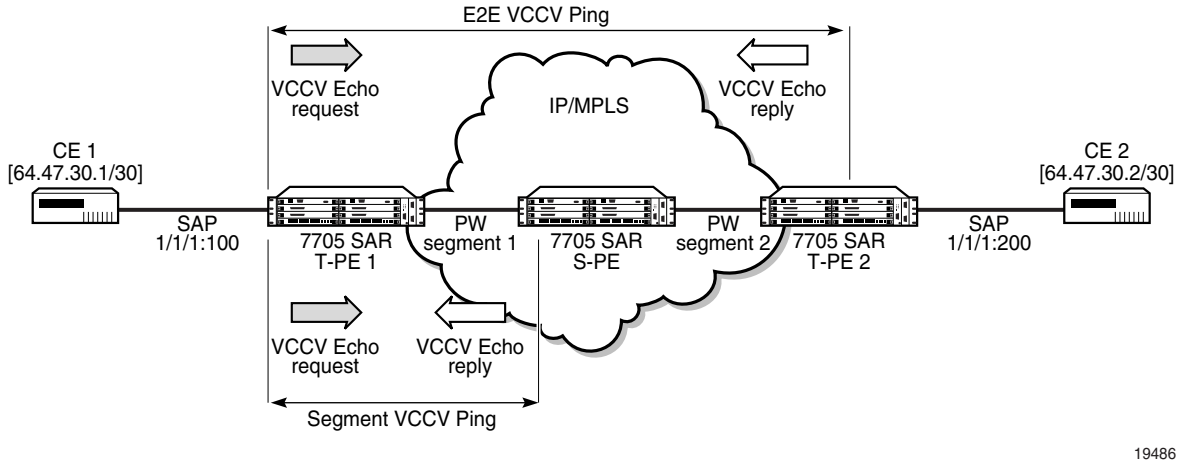

In the network, a Termination PE (T-PE) is where the pseudowire originates and terminates. The Switching PE (S-PE) is the node that performs pseudowire switching by cross-connecting two spoke SDPs.

VCCV ping on the 7705 SAR is capable of performing 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 7705 SAR 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 7705 SAR S-PE1 node does not process the VCCV OAM control word unless the VC label TTL expires. If the VC label TTL expires, the message is sent to the CSM for further validation and processing. This is the method described in *draft-hart-pwe3-segmented-pw-vccv*.

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. When an S-PE node receives a VCCV ping echo request destined for itself, it sends an IP-routed reply. VCCV trace can 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, where T-PE1 sends successive VCCV ping messages while incrementing the TTL value, starting from TTL=1. The procedure for each iteration is the same. 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 sent from the T-PE2 node or when a timeout occurs.

Automated VCCV Trace Capability for Multi-Segment Pseudowire

Although tracing of the MS-PW path is possible using the methods explained in the VCCV Ping section, these require multiple manual iterations and that require the FEC of the last pseudowire segment to the target T-PE/S-PE already be known 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 TTL values, starting from TTL=1.

The method is described in *draft-hart-pwe3-segmented-pw-vccv*, *VCCV Extensions for Segmented Pseudo-Wire*, and is pending acceptance by the PWE3 working group. In each iteration, the source T-PE or S-PE builds the MPLS echo request message in a way similar to VCCV Ping. The first message with TTL=1 will have the next-hop S-PE T-LDP session source address in the Remote PE Address field in the pseudowire FEC TLV. Each S-PE that terminates and processes the message will include the FEC 128 TLV corresponding to the pseudowire segment to its downstream node, in the MPLS echo reply message. The inclusion of the FEC TLV in the echo reply message is allowed according to *RFC 4379, Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. The source T-PE or S-PE then sends the next echo reply message with TTL=2 to test the next-next hop for the MS-PW. It copies the FEC TLV it received in the echo reply message into the new echo request message. The process is terminated when the reply is sent from the egress T-PE node 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 of VCCV trace can be displayed for a fewer number of pseudowire 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 the `min-ttl` value in order to correctly build the FEC of the desired subset of segments.

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-PW is supported with a mix of static and signaled pseudowire segments. However, VCCV ping and VCCV trace are not allowed if any segment of the MS-PW is static. Users cannot test a static segment or contiguous signaled segments of the MS-PW. VCCV ping and VCCV trace are not supported in static-to-dynamic configurations.

VCCV for MS-PW and Pseudowire Redundancy

VCCV is supported on S-PE nodes configured for MS-PW and PW redundancy. In this case, S-PE terminates the in-band or out-of-band (IP-routed) VCCV ping (echo reply) and can generate VCCV ping (echo request) toward the dynamic section of the PW segment.

To configure an S-PE for MS-PW and pseudowire redundancy, an explicit endpoint is required to configure the service. Only one explicit endpoint is supported. The first PW segment must be configured with a static inner label under an implicit endpoint. The second PW segment can be created as either a redundant or non-redundant PW using the explicit endpoint.

Note: A VLL service is in MS-PW and PW redundancy mode as long as there is one PW segment with an explicit endpoint configured.

On S-PE nodes configured for MS-PW and PW redundancy, each segment of the PW can be configured with its own independent control word. The control word of the dynamic segment does not have to match the control word of the static segment for traffic to flow. The control word is automatically inserted or removed from the packets as they are switched from one segment to the other based on the control word configuration for each segment.

From an OAM diagnostic perspective, only Type-1 VCCV is supported for the dynamic MS-PW segment, which means that the PW segment must be configured with the control word option. In this mode, the ability to support VCCVs is signaled through the label message and the optional VCCV TLV toward the dynamic segment on the S-PE. The S-PE terminates all VCCV packets arriving on the dynamic segment, then extracts them towards the CSM.

Detailed VCCV Trace Operation

In Figure 4, a trace can be performed on the MS-PW originating from T-PE1 by a single operational command. The following process occurs:

1. T-PE 1 sends a VCCV echo request with TTL set to 1 and a FEC 128 containing the pseudowire information of the first segment (pseudowire1 between T-PE 1 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-PE 2) and sends the echo reply back to T-PE 1.
3. T-PE 1 builds a second VCCV echo request based on the FEC 128 in the echo reply from the S-PE. It increments the TTL and sends the next echo request out to T-PE 2. 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-PE 2 receives and validates the echo request with the FEC 128 of the pseudowire2 from TPE 1. Since T-PE 2 is the destination node or the egress node of the MS-PW, 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-PE 1 receives the echo reply from T-PE 2. T-PE 1 is made aware that T-PE 2 is the destination of the MS-PW because the echo reply does not contain the FEC 128 and because its return code is 3. The trace process is completed.

VCCV Trace

VCCV trace is similar to LSP trace. VCCV trace is used to trace the entire path of a pseudowire (PW) with a single command.

VCCV trace is useful in multi-segment PW (MS-PW) applications where a single PW traverses one or more switched PEs (S-PEs). VCCV trace is an iterative process by which the initiating terminating PE (T-PE) sends successive VCCV ping messages, each message having an incrementing TTL value, starting from TTL=1. The procedure for each iteration is the same as that for VCCV-ping, where each node in which the VC label TTL expires will check the FEC and reply with the FEC to the downstream S-PE or far-end T-PE. The process is terminated when the reply is from the far-end T-PE or when a timeout occurs.

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

VPLS MAC Diagnostics

Although the LSP ping, SDP ping, and service ping tools enable transport tunnel testing and verify that 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 possible that even though 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. The 7705 SAR provides 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 — 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 — 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 a successful OAM and SAA test when there is a reply from a far-end node indicating the destination MAC address on an egress SAP or the CSM.
- CPE ping — 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 from 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. When 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, and so on. 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 operates like an LSP trace with 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.

When a MAC trace request is 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 as a unicast transmission 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 assigned by the system, where the source UDP port is really the demultiplexer that identifies the particular instance that sent the request, when that demultiplexer correlates the reply. The source IP address is the system IP address 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 address of the sender.

The Reply Mode is either 3 (control plane reply) or 4 (data plane reply), 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 and can detect 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 7705 SAR. Operators are encouraged to use the source IP address of 0.0.0.0 in order to prevent the provider's IP address from 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 as an OAM-induced learning with some other binding), to prevent new dynamic learning to over-write the existing OAM MAC entry, or 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, to 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 an operator to perform 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 request.

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 rather the control plane behavior of it.

Ethernet OAM Capabilities

The 7705 SAR supports Ethernet OAM capabilities, as described in the following sections:

- Ethernet OAM Overview
- 802.1ag and Y.1731 Functional Comparison
- ETH-CFM Ethernet OAM Tests (802.1ag and Y.1731)
- ITU-T Y.1731 Performance Monitoring (PM)
- EFM OAM (802.3ah)

Ethernet OAM Overview

The 7705 SAR supports the following Ethernet OAM capabilities:

- Ethernet Connectivity Fault Management (ETH-CFM) — for network layer OAM according to IEEE 802.1ag (dot1ag) and ITU Y.1731 standards, including loopbacks (LB), linktrace (LT), continuity checks (CC), and remote defect indicators (RDI). "Network layer" refers to an end-to-end context across a network.

  ITU-T Y.1731 provides functional enhancements to 802.1ag ETH-CFM, including alarm indication signals (AIS) and Ethernet signal tests (ETH-Test).

  See ETH-CFM Ethernet OAM Tests (802.1ag and Y.1731).
- Performance Monitoring (PM) — PM according to the ITU-T Y.1731 standard, including delay measurements (DM), delay variation measurements (DV), and loss measurements (LM).

  See ITU-T Y.1731 Performance Monitoring (PM).
- Ethernet First Mile (EFM) OAM — for the transport layer OAM according to IEEE 802.3ah (dot3ah) standards. "Transport layer" refers to a point-to-point link context or transport hop.

  See EFM OAM (802.3ah).

Ethernet OAM capabilities on the 7705 SAR are similar to the OAM capabilities offered in SONET/SDH networks and include loopback tests to verify end-to-end connectivity, test pattern generation (and response) to verify error-free operation, and alarm message generation in case of fault conditions to ensure that the far end is notified of the failure.

Ethernet OAM configurations are maintained across Control and Switching module (CSM) switchovers.

Ethernet OAM Usage Examples

Figure 5 illustrates the complementary use of dot3ah and dot1ag to locate points of failure along a route from BTS to BSC. Since dot1ag and Y.1731 have similar functions, only dot1ag is discussed in order to simplify the explanation.

In Figure 5, from the IP/MPLS (network) layer perspective, the 7705 SAR looks as though it is connected directly to the 7750 SR. From the Ethernet (transport) layer perspective, the route passes through many ports and nodes, where each port or node is a potential point of failure. These failure points cannot be detected using IP/MPLS OAM capabilities (that is, using ETH-CFM (dot1ag)). However, they can be detected using EFM OAM (dot3ah) capabilities.

Figure 5: 7705 SAR Ethernet OAM Endpoints

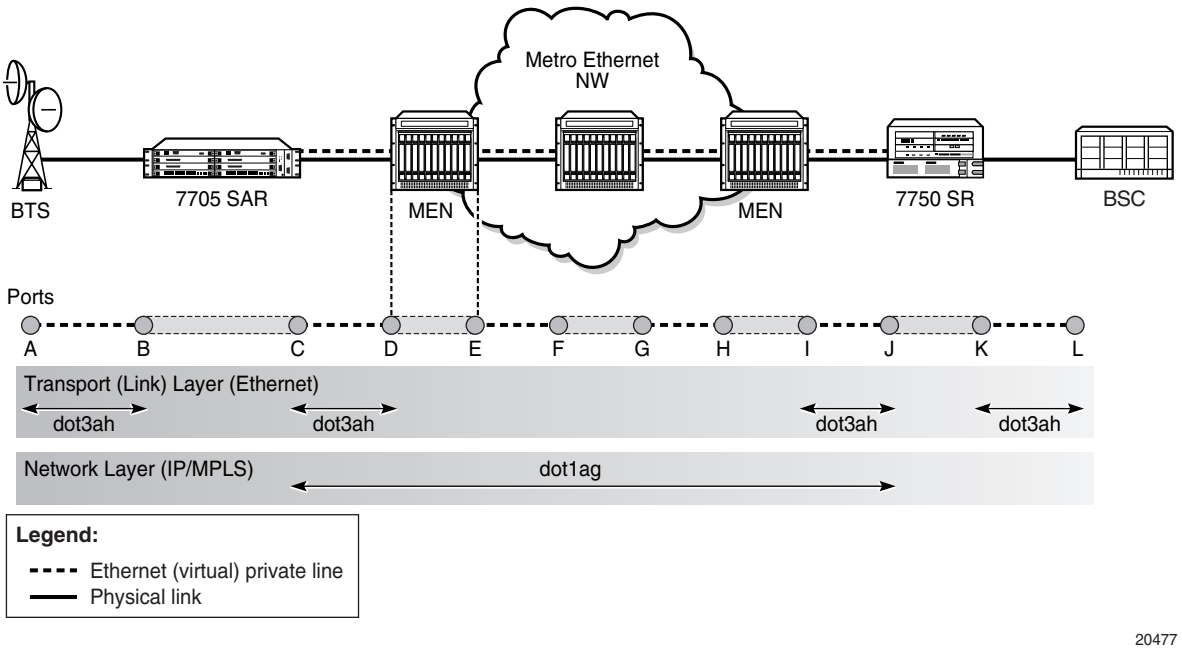

Dot3ah uses port-level loopbacks to check and verify last-mile Ethernet frame integrity, connectivity verification between ports and nodes, and so on. As shown in Figure 5, dot3ah provides transport (link) layer OAM between the BTS and the 7705 SAR access port facing the BTS (ports A and B), or between the 7705 SAR network port and the MEN switch (ports C and D). Ethernet first mile (EFM) OAM allows users to test frame integrity and detect Ethernet layer failures faster than using associated heart-beat messages.

Dot1ag checks end-to-end connectivity across an Ethernet PW (across a network). Since end-to-end connectivity differs depending on the service provided and the span of the network, dot1ag can operate at several MD levels (as defined in the IEEE 802.1ag standard). For example, in Figure 5, ETH-CFM (dot1ag) could be used by a MEN provider at one MD level to ensure connectivity between ports D and I (or possibly all the way to their customer's Ethernet ports, C and J). Similarly, a mobile backhaul service provider (MBSP) can use dot1ag at another MD level to ensure connectivity between ports B and K (and possibly between ports A and L).

Figure 6 and Figure 7 illustrate the use of ETH-CFM to verify connectivity across an Ethernet PW and EFM OAM to verify transport layer connectivity between two directly connected nodes.

For example, in Figure 6, an MBSP can use dot1ag between the two Ethernet spoke SDP endpoints (ports C and J, which define the Ethernet PW) to ensure connectivity. Similarly, a MEP can use dot1ag between ports D and I to ensure the health status of the Ethernet (virtual) private line.

Figure 6: ETH-CFM (Dot1ag) Capabilities on the 7705 SAR

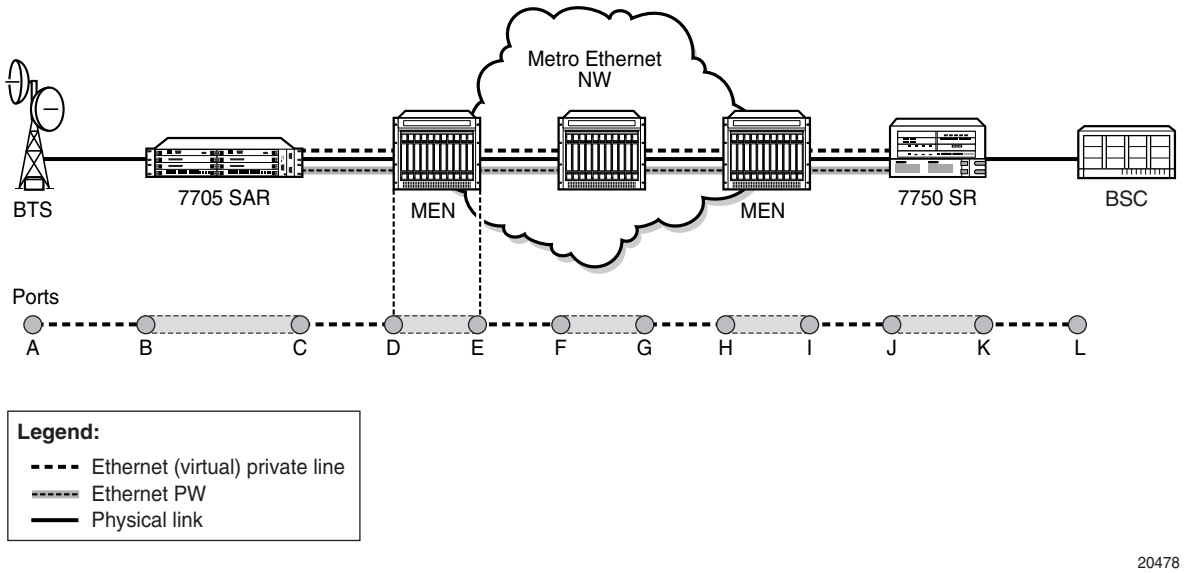

In Figure 7, EFM OAM ensures transport layer connectivity between two directly connected nodes. Figure 7 illustrates three scenarios in which EFM can be used by the MEN provider to ensure error-free connectivity to the 7705 SAR (the cell site) via loopback tests, including:

- scenario 1: EFM termination at the Ethernet access port, which includes loopback tests, heart-beat messages at the Ethernet layer with dying gasp, and termination of customer device-initiated EFM packets at the access port
- scenario 2: EFM termination at the Ethernet network port, which includes network-side loopbacks
- scenario 3: EFM tunneling through an Epipe service

Figure 7: EFM OAM (Dot3ah) Capabilities on the 7705 SAR

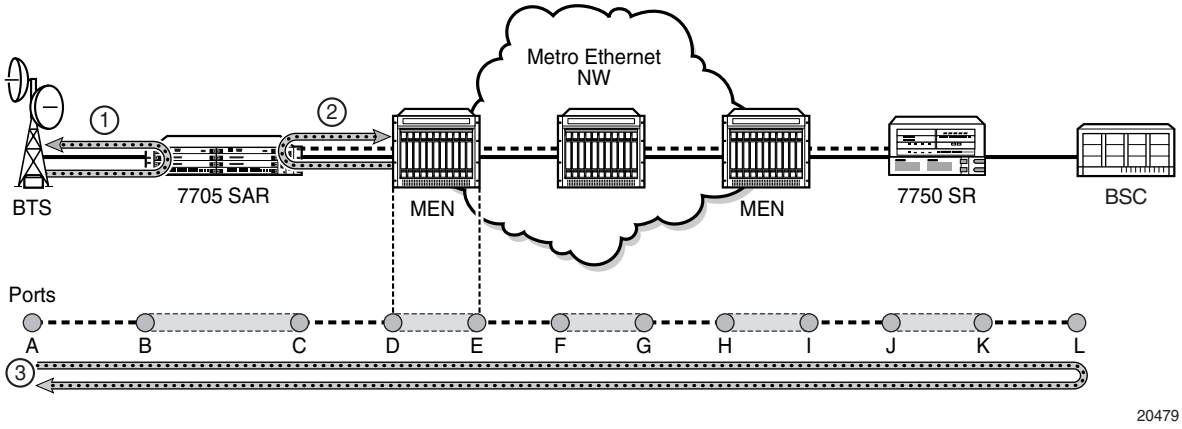

802.1ag and Y.1731 Functional Comparison

Table 3 lists the 802.1ag and Y.1731 OAM functions supported on the 7705 SAR. For each function and test, the table identifies the PDU that carries the test data, the test's target entity, and the standard(s) that the 7705 SAR supports for the test.

For example, the 7705 SAR can run an Ethernet Continuity Check using an ETH-CC test according to the dot1ag and the Y.1731 standards. For either standard, the test data is carried in a Continuity Check message (CCM) and the test target is a MEP.

The Fault Management (FM) capabilities of ITU-T Y.1731 extend the functionality of dot1ag (ETH-CFM) with additional FM functions as well as performance management (PM) capabilities. The generation of AIS and RDI messages are defined under the FM section of the Y.1731 specification, whereas Ethernet layer, delay, jitter, loss, and throughput tests are part of Y.1731 PM capabilities.

Table 3: 802.1ag and Y.1731 OAM Functionality Overview

Test	OAM Function	PDU	Target	Standard
ETH-LB	Loopback	LBM, LBR	MEP	dot1ag, Y.1731
ETH-LT	Linktrace	LTM, LTR	MEP	dot1ag, Y.1731
ETH-CC	Continuity Check	CCM	MEP	dot1ag, Y.1731
ETH-RDI	Remote Defect Indication	CCM	MEP	dot1ag, Y.1731
ETH-AIS	Alarm Indication Signal	AIS	MEP	Y.1731
ETH-LM	Frame Loss Measurement (dual-ended)	CCM	MEP	Y.1731
ETH-LM	Frame Loss Measurement (single-ended)	LMM, LMR	MEP	Y.1731
ETH-DM	Frame Delay Measurement (two-way)	DMM, DMR	MEP	Y.1731
ETH-DM	Frame Delay Measurement (one-way)	1DM	MEP	Y.1731
ETH-DV	Frame Delay Variation (one-way)	DMM, DMR	MEP	Y.1731
ETH-Test	Test Error Measurements	TST	MEP	Y.1731

ETH-CFM Ethernet OAM Tests (802.1ag and Y.1731)

Ethernet Connectivity Fault Management (ETH-CFM) is defined in the IEEE 802.1ag and ITU Y.1731 standards. It specifies protocols, procedures, and managed objects to support fault management (including discovery and verification of the path), detection, and isolation of a connectivity fault for each Ethernet service instance.

IEEE 802.1ag and Y.1731 can detect:

- loss of connectivity
- unidirectional loss
- loops
- merging of services

The implementation of Y.1731 on the 7705 SAR also provides the following enhancements:

- Ethernet Alarm Indication Signal (ETH-AIS)
- Ethernet Test function (ETH-Test)

ETH-CFM uses Ethernet frames and can be distinguished by its Ethertype value (8902). With ETH-CFM, interoperability can be achieved between different vendor equipment in the service provider network, up to and including customer premises bridges.

ETH-CFM is configured at both the global level and the Ethernet service level. The following entities and their configuration levels are listed below:

- global level
  - → MA and MEG
  - → MD
  - → MD level and MEG level
- Ethernet service level
  - → MEP

For information on configuring ETH-CFM to set up an Ethernet OAM architecture, refer to the “ETH-CFM (802.1ag and Y.1731)” section in the 7705 SAR OS Services Guide.

Loopback (LB)

The loopback function is supported by 802.1ag and Y.1731 on the 7705 SAR. A Loopback Message (LBM) is generated by a MEP to its peer MEP. Both dot1ag and dot3ah loopbacks are supported. The loopback function is similar to IP or MPLS ping in that it verifies Ethernet connectivity between the nodes on a per-request basis. That is, it is non-periodic and is only initiated by a user request.

In Figure 8, the line labeled LB represents the dot1ag loopback message between the 7750 SR (source) and 7705 SAR (target) over an Epipe. The 7750 SR-generated LBM is switched to the 7705 SAR, where the LBM message is processed. Once the 7705 SAR generates the Loopback Reply message (LBR), the LBR is switched over the Ethernet PW to the 7750 SR.

Figure 8: Dot1ag Loopback Test

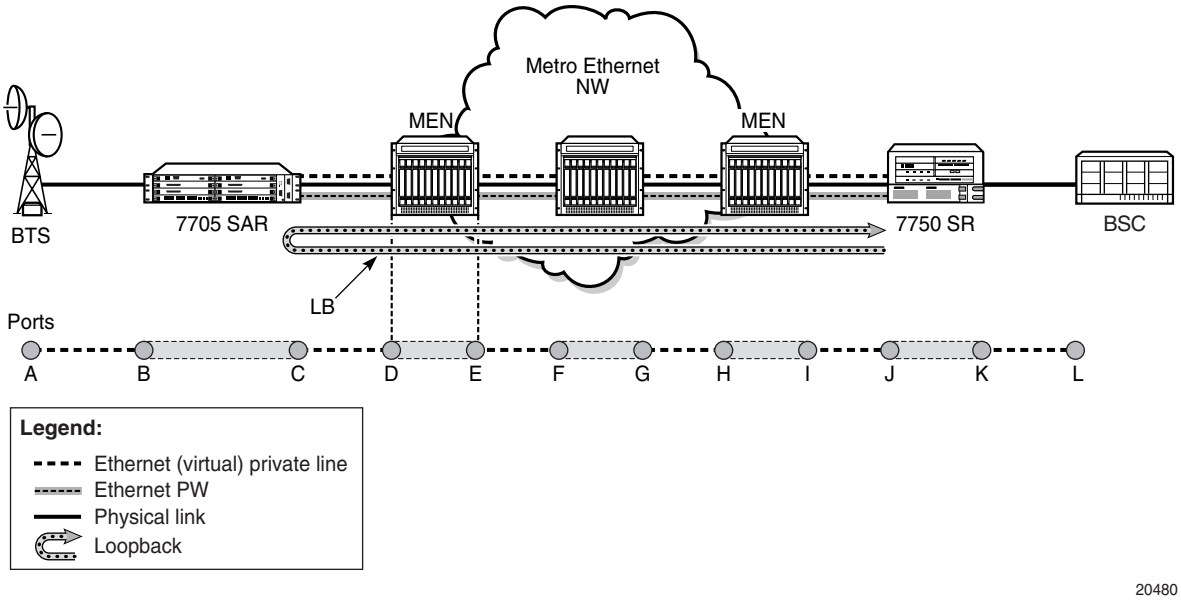

Linktrace (LT)

The linktrace function is supported by 802.1ag and Y.1731 on the 7705 SAR. A Linktrace Message (LTM) is originated by a MEP and targeted to a peer MEP in the same MA and within the same MD level. Its function is similar to IP traceroute. The peer MEP responds with a Linktrace Reply (LTR) message after successful inspection of the LTM.

Throughput Measurement

Throughput measurement is performed by sending frames to the far end at an increasing rate (up to wire speed) and measuring the percentage of frames received back. In general, the rate is dependent on frame size; the larger the frame size, the lower the rate.

The Y.1731 specification recommends the use of unicast ETH-LB and ETH-Test frames to measure throughput.

On the 7705 SAR, LBM processing and LBR generation is enhanced and occurs on the datapath, allowing the node to respond to loopback messages at wire speed and making in-service throughput tests possible. Thus, if the 7705 SAR receives LBMs at up to wire speed, it can generate up to an equal number of LBRs. In order to process LBMs at wire speed, there must be either no TLVs or a single TLV (which is a Data TLV) in the LBM frame. The End TLV field (0) must be present and the frame can be padded with data after the End TLV field in order to increase the size of the frame. The MAC address cannot be a multicast MAC address; it must be the MEP MAC destination address (DA).

Datapath processing of LBMs is supported for the following MEPs:

- dot1ag
  - → SAP Up MEP
  - → SAP Down MEP
  - → spoke-SDP Down MEP
- Y.1731
  - → SAP Up MEP
  - → SAP Down MEP

For spoke-SDP Down MEPs, fastpath (datapath) LBM processing requires that both interfaces—the LBM receiver and the LBR transmitter—reside on the same adapter card. For example, if the 7705 SAR must perform a reroute operation and needs to move the next hop interface to another adapter card (that is, LBMs are received on one card and LBRs are transmitted on another), then the fastpath processing of LBMs is terminated and LBM processing continues via the CSM.

Continuity Check (CC)

The continuity check function is supported by 802.1ag and Y.1731 on the 7705 SAR. A Continuity Check Message (CCM) is a multicast frame that is generated by a MEP and sent to its remote MEPs in the same MA. The CCM does not require a reply message. To identify faults, the receiving MEP maintains a MEP database with the MAC addresses of the remote MEPs with which it expects to maintain connectivity checking. The MEP database can be provisioned manually. If there is no CCM from a monitored remote MEP in a preconfigured period, the local MEP raises an alarm.

The following CC capabilities are supported:

- enable and disable CC for a MEP
- automatically put local MEPs into the database when they are created
- manually configure and delete the MEP entries in the CC MEP monitoring database. Note that the only local provisioning required to identify a remote MEP is the remote MEP identifier (using the `remote-mepid` *`mep-id`* command).
- CCM transmit interval: 10ms, 100ms, 1s, 10s, 1m, 10m (default: 10s)
- transmit interval: 10ms, 100ms, 1s, 10s, 1m, 10m (default: 10s)
- CCM declares a fault when it:
  - → stops hearing from one of the remote MEPs for a period of 3.5 times the CC interval
  - → hears from a MEP with a lower MD level
  - → hears from a MEP that is not in the same MA
  - → hears from a MEP that is in the same MA but is not in the configured MEP list
  - → hears from a MEP that is in the same MA with the same MEP ID as the receiving MEP
  - → recognizes that the CC interval of the remote MEP does not match the local configured CC interval
  - → recognizes that the remote MEP declares 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.
- CC must be enabled in order for RDI information to be carried in the CCM OAMPDU

ETH-RDI

The Ethernet Remote Defect Indication function (ETH-RDI) is used by a MEP to communicate to its peer MEPs that a defect condition has been encountered. Defect conditions such as signal fail and AIS may result in the transmission of frames with ETH-RDI information. ETH-RDI is used only when ETH-CC transmission is enabled.

ETH-RDI has the following two applications:

- single-ended fault management — the receiving MEP detects an RDI defect condition, which gets correlated with other defect conditions in this MEP. The absence of received ETH-RDI information in a single MEP indicates the absence of defects in the entire MEG.
- contribution to far-end performance monitoring — the transmitting MEP reflects that there was a defect at the far end, which is used as an input to the performance monitoring process

A MEP that is in a defect condition transmits frames with ETH-RDI information. A MEP, upon receiving frames with ETH-RDI information, determines that its peer MEP has encountered a defect condition.

The specific configuration information required by a MEP to support the ETH-RDI function is as follows:

- MEG level — the MEG level at which the MEP exists
- ETH-RDI transmission period — application-dependent and is the same value as the ETH-CC transmission period
- priority — the priority of frames containing ETH-RDI information and is the same value as the ETH-CC priority

The PDU used to carry ETH-RDI information is the CCM.

Note: There is a scenario in which an RDI defect is transmitted when it should not be transmitted. Normally, to indicate a failure, an Up MEP transmits a Port or Interface Status TLV (or both). However, an Up MEP will also transmit ETH-RDI if the port or interface on which it resides experiences a failure. In addition, the normal expectation is for the MEP to report no defects. However, the MEP will instead indicate a DefRemoteCCM defect.

Furthermore, under normal conditions, the remote MEP would indicate only a DefMACstatus. However, the remote MEP indicates both a DefMACstatus and a DefRDICCM defect.

ETH-AIS

The Ethernet Alarm Indication Signal function (ETH-AIS) is a Y.1731 CFM enhancement used to suppress alarms at the client (sub) layer following detection of defect conditions at the server (sub) layer.

Transmission of frames with ETH-AIS information can be enabled or disabled on a Y.1731 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 where ETH-CC is enabled
- AIS condition in the case where ETH-CC is disabled

For a point-to-point Ethernet 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 a defect condition upon receiving a frame with ETH-AIS information. Alarm suppression is simplified by the fact that a MEP is expected to suppress only those defect conditions associated with its peer MEP.

Only a MEP, including a server MEP, is configured to issue frames with ETH-AIS information. Upon detecting a defect condition, the MEP can immediately start transmitting periodic frames with ETH-AIS information at a configured client MEG level. A MEP continues to transmit periodic frames with ETH-AIS information until the defect condition is removed. Upon receiving a frame with ETH-AIS information from its server (sub) layer, a client (sub) layer MEP detects the AIS condition and suppresses alarms associated with all its peer MEPs. Once the AIS condition is cleared, a MEP resumes alarm generation upon detecting defect conditions.

The following specific configuration information is required by a MEP to support ETH-AIS:

- client MEG level — the MEG level at which the most immediate client layer MEPs exist
- ETH-AIS transmission period — the transmission period of frames with ETH-AIS information
- priority — the priority of frames with ETH-AIS information

ETH-Test

The Ethernet Test (signal) function (ETH-Test) is a Y.1731 CFM enhancement used to perform one-way, on-demand, in-service diagnostics tests, which include verifying frame loss, bit errors, and so on.

Note: The out-of-service diagnostics test is not supported in the 7705 SAR.

When configured to perform such tests, a MEP inserts frames with ETH-Test information such as frame size and transmission patterns.

When an in-service ETH-Test function is performed, data traffic is not disrupted and the frames with ETH-Test information are transmitted.

To support ETH-Test, a Y.1731 MEP requires the following configuration information:

- MEG level — the MEG level at which the MEP exists
- unicast MAC address — the unicast MAC address of the peer MEP for which the ETH-Test is intended
- data — an optional element with which to configure data length and contents for the MEP. The contents can be a test pattern and an optional checksum.

  Examples of test patterns include all 0s or all 1s. At the transmitting MEP, this configuration information is required for a test signal generator that is associated with the MEP. At the receiving MEP, this configuration is required for a test signal detector that is associated with the MEP.
- priority — the priority of frames with ETH-Test information

A MEP inserts frames with ETH-Test information towards a targeted peer MEP. The receiving MEP detects the frames with ETH-Test information and performs the requested measurements.

ITU-T Y.1731 Performance Monitoring (PM)

The Y.1731 Performance Monitoring (PM) functions can be used to measure Ethernet frame delay, delay variation, throughput (including throughput at queue-rates), and frame loss. These performance parameters are defined for point-to-point Ethernet connections.

Delay and Delay Variation Measurements (DM and DV)

The Y.1731 recommendation covers the following performance parameters, which are based on Metro Ethernet Forum (MEF) 10:

- frame delay — specified as one-way or round-trip delay for a frame, where frame delay is defined as the time elapsed since the start of transmission of the first bit of the frame by a source node until the reception of the frame by the destination node or the same source node
- frame delay variation — a measure of the variations in the frame delay between a pair of service frames, where the service frames belong to the same CoS instance on a point-to-point Ethernet connection

The performance parameters listed above are applicable to Ethernet services frames. Services frames are those frames that conform to an agreed-upon level of bandwidth profile conformance and are associated with a particular CoS identifier. Services frames are admitted at the ingress Ethernet flow point of a point-to-point Ethernet connection and should be delivered to the egress Ethernet flow point.

The 7705 SAR supports one-way and two-way Ethernet Delay Measurement (ETH-DM) (section 8.2 of the Y.1731 standard), using the CLI commands `oam>eth-cfm>one-way-delay-test` and `config>service>epipe>two-way-delay-test`. Ethernet Delay Variation measurements (ETH-DV) are run along with the one-way and two-way ETH-DM tests.

For ETH-DM, the accuracy of the measurement is in the microseconds range.

Y.1731 Delay Measurement (DM)

Y.1731 delay measurement implementation ensures the most accurate results under all circumstances. The implementation ensures that there is minimal delay measurement error between packet generation and packet play-out over the Ethernet link.

In order to isolate delay measurement results from the effects of any queuing, scheduling, and shaping procedures, timestamping of DM frames in the transmit direction is performed when the first byte of the DM frame is put on the wire (that is, once the actual serialization has started). Last-minute timestamping ensures that DM tests truly measure the delay between two SAP or port endpoints, and not the delay imposed by the routers. Using these accurate measurements, a network operator can separate the delay induced by the routers from the transmission delay introduced by the transmission network, such as a Metro Ethernet network (MEN) or Generic Framing Procedure (GFP) over SONET links.

Timestamping of DM frames in the receive direction is similar to last-minute timestamping in the transmit direction, except that the timestamp on received DM frames occurs when the last byte is received from the wire. Last-minute timestamping ensures the ability to separate the total delay measurement into a node component and a transport network component.

Last-minute timestamping is used for both one-way and two-way delay test frames to ensure accuracy.

Loss Measurement (LM)

The 7705 SAR supports single-ended and dual-ended Ethernet Loss Measurement (ETH-LM) tests. Dual-ended LM tests are enabled under the `config>service>epipe>sap>eth-cfm>mep` context. Once enabled, dual-ended LM tests run continuously in the background. Single-ended LM tests are run from the `oam>eth-cfm` context and are considered on-demand tests.

Y.1731 loss measurement functionality is implemented to ensure the most accurate results under all circumstances. Each adapter card has a network processor (NP). LM counters are maintained at the NP. The NP is responsible for incrementing and resetting these counters. These counters are accessed by the CSM CPU in order to calculate and display the loss (percentage) to the user.

LM/CCM frames follow the associated QoS path and therefore might inadvertently report loss due to local congestion even before the frame is switched onto the link. In order to reflect the true experience of a particular QoS setting, generated LM/CCM frames follow the egress QoS path. Once generated, these frames are classified in the same manner as the applicable dot1p-to-FC mapping, associated queuing, and scheduling rules. Following the proper path ensures that loss measurements reflect the experience of a given FC all the way through the network, including within the 7705 SAR platform. As is the case for any other frame of the same FC (that is, user or control frame), the LM/CCM frame follows the associated QoS path to reflect the real experience.

For example, newly generated LM/CCM frames that have a higher counter value can be forwarded sooner than LM/CCM frames with a lower counter value that have been generated but are waiting to be serviced (that is, frames with a lower queue, a queue in the out-of-profile state; or a single SAP with multiple FCs). As a result, when under congestion, the LM ratio would increase to reflect local loss if lower-priority frames cannot be serviced in a timely manner.

In addition, congestion, and hence prioritization, can occur anywhere in the transport network, which means that a reordering might take place not only on the ingress point, but anywhere in the network along the entire path.

The loss ratio is calculated based on the aggregate frames being transmitted and received. Thus, in a uncongested network, the loss ratio would be 0%. With congestion, not all frames may be sent out to the network (that is, higher priority traffic, and so on) or any one of the transit nodes or the endpoint node might drop the packet, which would end up with loss.

The above-described behavior for following the QoS path equally applies to both Up and Down MEPs. Loss measurements in both up and down directions for the same MEP can be performed simultaneously.

The counters used for loss measurement in LM and CCM frames are appended as late as possible in the datapath. Appending the counters at the last minute to the LM or CCM frames ensures that a scheduling priority issue or some other queue-delaying event does not delay the OAM frame in a queue. If the counters are updated or generated earlier in the datapath, then the OAM frames could be affected by queuing or scheduling delays, which might cause the frames to be counted as lost frames when the far-end receive timer expires.

The following notes apply to Y.1731 LM tests.

- Single-ended and dual-ended LM tests cannot be enabled on a MEP simultaneously. That is, either a single-ended or a dual-ended LM test can be enabled on a given MEP at any given time.
- The behavior and the interaction between single- and dual-ended LM tests are described in the following list. Error conditions, such as correct domain level and valid destination address (DA) MAC, are not covered in the list:
  - → if dual-ended loss test is disabled:
    - CCM frames are transmitted with LM counters set to 0
    - CCM frames being received are not processed for LM
    - LMM and LMR frames being received are processed
    - single-ended tests can be enabled (not blocked by CLI)
  - → if dual-ended loss test is enabled:
    - single-ended tests cannot be enabled (blocked by CLI)
    - LMM and LMR frames being received will be dropped

- Multiple MEPs bound to the same Epipe SAP but belonging to different MEG levels can perform LM tests simultaneously.
- CCM must be enabled before a dual-ended LM test can be enabled.
- When a dual-ended LM test is enabled, the user cannot disable CCM. The dual-ended LM test needs to be disabled before the CCM can be disabled.
- For dual-ended LM tests, an alarm is declared when frame losses are greater than an alarm threshold configured for the MEP. The granularity of the alarm threshold (declaring or clearing) is 0.01%. The default threshold is set to 0.25%.
- On a per-SAP basis, there is one set of Rx and Tx Local LM counters and one set of Rx and Tx Remote LM counters. LM counters are not separated on a per-MAC source address (SA) basis. All MEPs, irrespective of their MD or MEG level, share the same set of Rx and Tx LM counters.
- On the CLI, there are interval counters and accumulated counters. The CCM counters are referred to as Local and FarEnd counters and the accumulated counters are referred to as Near-End and Far-End counters.
- The LM counters are incremented when a user data frame reaches a SAP. Since there is only one set of Tx and Rx Local counters per SAP, each user data frame received by all the MEPs configured on that SAP is counted.
- OAM frames with MEP levels matching or lower than the locally configured MEP level are not counted. They are treated and processed as OAM frames. This functionality applies to both received and transmitted OAM frames. CFM OAM frames at higher MEP levels are counted as user data frames.
  - → For example, assume a SAP with two MEPs configured on it; one MEP at level 5 and the other at level 6.
  - → When a level 6 OAM frame is received, it is extracted to the CSM for processing and is not counted by LM counters. It is treated as an OAM frame.
  - → The same behavior applies in the transmit direction. In the above example, any level 5 or level 6 OAM frames generated by the local SAR would not be counted by the far-end LM counters.
- For dual-ended LM tests, any received CCMs with all LM counters being 0s (zeros) are treated as invalid. In this case, the 7705 SAR resets the LM counters for the current and previous CCMs to 0s (zeros). Accumulated counters are not reset.
- Except for a valid counter rollover scenario, if the value of any CCM/LMR counter is less than the value of the same counter in the previous CCM/LMR frame, then the accumulated values of all counters are not increased; they are kept at the same values as before the last CCM/LMR frame is processed.
- When the first valid CCM/LMR frame — that is, a frame with at least one non-zero LM counter — is received after a dual-ended loss test is enabled or a single-ended loss test is launched, the accumulated values cannot be calculated. In this case, the counters are resaved as current counters. When the next received CCM/LMR frame with valid LM counts is received, it will trigger the update of accumulation counts.

Accumulated counts always start at 0 for each launch of a single-ended test. However, the accumulation counts do not change nor do they get reset to all 0s when dual-ended loss tests become disabled. For dual-ended loss tests, accumulated counts can be restarted at 0s by removing the existing LM result of a particular MEP with the CLI command `clear>eth-cfm>dual-ended-loss-test>mep` *mep-id* `domain` *md-index* `association` *ma-index*, or the equivalent SNMP command.

EFM OAM (802.3ah)

802.3ah Clause 57 defines the Ethernet in the First Mile (EFM) OAM sublayer, which is a link level Ethernet OAM that is supported on 7705 SAR Ethernet ports configured as network or access ports. It provides mechanisms for monitoring link operations such as remote fault indication and remote loopback control. Ethernet OAM gives network operators the ability to monitor the health of Ethernet links and quickly determine the location of failing links or fault conditions.

Because some of the sites where the 7705 SAR will be deployed will have only Ethernet uplinks, this OAM functionality is mandatory. For example, mobile operators must be able to request remote loopbacks from the peer router at the Ethernet layer in order to debug any connectivity issues. EFM OAM provides this capability.

EFM OAM defines a set of events that may impact link operation. The following events are supported:

- critical link events (defined in 802.3ah clause 57.2.10.1)
  - → link fault: the PHY has determined that a fault has occurred in the receive direction of the local DTE
  - → dying gasp: an unrecoverable local failure condition has occurred
  - → critical event: an unspecified critical event has occurred

These critical link events are signaled to the remote DTE by the flag field in OAMPDUs.

EFM OAM is supported on network and access Ethernet ports, and is configured at the Ethernet port level. The access ports can be configured to tunnel the OAM traffic originated by the far-end devices.

EFM OAM has the following characteristics.

- All EFM OAM, including loopbacks, operate on point-to-point links only.
- EFM loopbacks are always line loopbacks (line Rx to line Tx).
- When a port is in loopback, all frames (except EFM frames) are discarded. If dynamic signaling and routing is used (dynamic LSPs, OSPF, IS-IS, or BGP routing), all services also go down. If all signaling and routing protocols are static (static routes, LSPs, and service labels), the frames are discarded but services stay up.

The following EFM OAM functions are supported:

- OAM capability discovery
- configurable transmit interval with an Information OAMPDU
- active or passive mode
- OAM loopback
- OAMPDU tunneling and termination (for Epipe service)
- dying gasp at network and access ports

For information on Epipe service, refer to the 7705 SAR OS Services Guide, “Ethernet VLL (Epipe) Services”.

Unidirectional OAM Operation

Some physical layer devices support unidirectional OAM operation. When a link is operating in unidirectional OAM mode, the OAM sublayer ensures that only information OAMPDUs with the Link Fault critical link event indication set and no Information TLVs are sent across the link.

Remote Loopback

EFM OAM provides a link-layer frame loopback mode, which can be controlled remotely.

To initiate a remote loopback, the local EFM OAM client enables the OAM remote loopback command to send a loopback control OAMPDU. After receiving the loopback control OAMPDU, the remote OAM client puts the remote port into local loopback mode.

OAMPDUs are slow protocol frames that contain appropriate control and status information used to monitor, test, and troubleshoot OAM-enabled links.

To exit a remote loopback, the local EFM OAM client sends a loopback control OAMPDU by disabling the OAM remote loopback command. After receiving the loopback control OAMPDU, the remote OAM client puts the port back into normal forwarding mode.

When a port is in local loopback mode (the far end requested an Ethernet OAM loopback), any packets received on the port will be looped back, except for EFM OAMPDUs. No data will be transmitted from the node; only data that is received on the node will be sent back out.

When the node is in remote loopback mode, local data from the CSM is transmitted, but any data received on the node is dropped, except for EFM OAMPDUs.

Remote loopbacks should be used with caution; if dynamic signaling and routing protocols are used, all services go down when a remote loopback is initiated. If only static signaling and routing is used, the services stay up. On the 7705 SAR, the Ethernet port can be configured to accept or reject the remote-loopback command.

802.3ah OAMPDU Tunneling and Termination for Epipe Services

Customers who subscribe to Epipe service might have customer equipment running 802.3ah at both ends. The 7705 SAR can be configured to tunnel EFM OAMPDUs received from a customer device to the other end through the existing network using MPLS or GRE, or to terminate received OAMPDUs at a network or an access Ethernet port.

Note: This feature applies only to port-based Epipe SAPs because 802.3ah runs at port level, not at VLAN level.

While tunneling offers the ability to terminate and process the OAM messages at the head-end, termination on the first access port at the cell site can be used to detect immediate failures or can be used to detect port failures in a timelier manner.

The user can choose either tunneling or termination, but not both at the same time.

In Figure 7, scenario 1 shows the termination of received EFM OAMPDUs from a customer device on an access port, while scenario 2 shows the same thing except for a network port. Scenario 3 shows tunneling of EFM OAMPDUs through the associated Ethernet PW. To configure termination (scenario 1), use the `config>port>ethernet>efm-oam>no shutdown` command.

Dying Gasp

Dying gasp is used to notify the far end that EFM-OAM is disabled or shut down on the local port. The dying gasp flag is set on the OAMPDUs that are sent to the peer. The far end can then take immediate action and inform upper layers that EFM-OAM is down on the port.

When a dying gasp is received from a peer, the node logs the event and generates an SNMP trap to notify the operator.

Ethernet Loopbacks

The 7705 SAR supports the following loopbacks on Ethernet ports:

- timed line loopbacks
- timed line loopbacks with MAC address swapping
- both timed and untimed internal loopbacks (equipment loopbacks)
- CFM loopbacks for OAM

Line and Internal Ethernet Loopbacks

A line loopback loops frames received on the corresponding port back towards the transmit direction. Line loopbacks are supported on ports configured in network mode.

Similarly, a line loopback with MAC addressing loops frames received on the corresponding port back towards the transmit direction, and swaps the source and destination MAC addresses before transmission. See MAC Swapping for more information.

An internal loopback loops frames from the local router back to the framer. This is usually referred to as an equipment loopback. The transmit signal is looped back and received by the interface. Internal loopbacks are supported on ports configured in access mode.

If a loopback is enabled on a port, the port mode cannot be changed until the loopback has been disabled.

A port can support only one loopback at a time. If a loopback exists on a port, it must be disabled or the timer must expire before another loopback can be configured on the same port. EFM-OAM cannot be enabled on a port that has an Ethernet loopback enabled on it. Similarly, an Ethernet loopback cannot be enabled on a port that has EFM-OAM enabled on it.

When an internal loopback is enabled on an Ethernet port, autonegotiation is turned off silently. This is to allow an internal loopback when the operational status of a port is down. Any user modification to autonegotiation on a port configured with an internal Ethernet loopback will not take effect until the loopback is disabled.

The loopback timer can be configured from 30 seconds to 86400 seconds. All Ethernet loopbacks are turned off automatically under the following conditions: an adapter card reset, an activity switch, or timer expiry. The timer for an internal loopback can also be configured to 0 seconds, turning it into a latched loopback that is enabled indefinitely, until it is turned off by the user or there is a system restart. These latched loopbacks survive adapter card resets and activity switches.

The `admin-save` and `admin-save-detail` commands do not save Ethernet loopbacks to the database.

MAC Swapping

Typically, an Ethernet port loopback only echoes back received frames. That is, the received source and destination MAC addresses are not swapped. However, not all Ethernet equipment supports echo mode, where the original sender of the frame must support receiving its own port MAC address as the destination MAC address.

The MAC swapping feature on the 7705 SAR is an optional feature that will swap the received destination MAC address with the source MAC address when an Ethernet port loopback is in line mode. After the swap, the FCS is recalculated to ensure the validity of the Ethernet frame and to ensure that the frame is not dropped by the original sender due to CRC error.

Interaction of Ethernet Port Loopback with Other Features

EFM OAM and line loopback are mutually exclusive. If one of these functions is enabled, it must be disabled before the other can be used.

However, a line loopback precedes the dot1x behavior. That is, if the port is already dot1x-authenticated it will remain so. If it is not, EAP authentication will fail.

CFM Loopbacks for OAM on Ethernet Ports

Connectivity fault management (CFM) loopback support for loopback messages (LBMs) on Ethernet ports allows operators to run standards-based Layer 1 and Layer 2 OAM tests on ports receiving unlabeled packets.

Prior to Release 4.0, the 7705 SAR supported CFM MEPs associated with different endpoints (that is, spoke SDP down MEPs, and SAP up and SAP down MEPs). In addition, for traffic received from an uplink (network ingress), the 7705 SAR supported CFM only for labeled packets. Release 4.0 adds CFM LBM support for unlabeled packets. CFM loopbacks are applied to the Ethernet port.

See Ethernet OAM Capabilities for information on CFM MEPs.

Figure 9 shows an application where an operator leases facilities from a transport network provider in order to transport traffic from a cell site to their MTSO. The operator leases a certain amount of bandwidth between the two endpoints (the cell site and the MTSO) from the transport provider, who offers Ethernet Virtual Private Line (EVPL) or Ethernet Private Line (EPL) PTP service. Before the operator offers services on the leased bandwidth, the operator runs OAM tests to verify the SLA. Typically, the transport provider (MEN provider) requires that the OAM tests be run in the direction of (towards) the first Ethernet port that is connected to the transport network. This is done in order to eliminate the potential effect of queuing, delay, and jitter that may be introduced by a spoke SDP or SAP.

Figure 9: CFM Loopback on Ethernet Ports

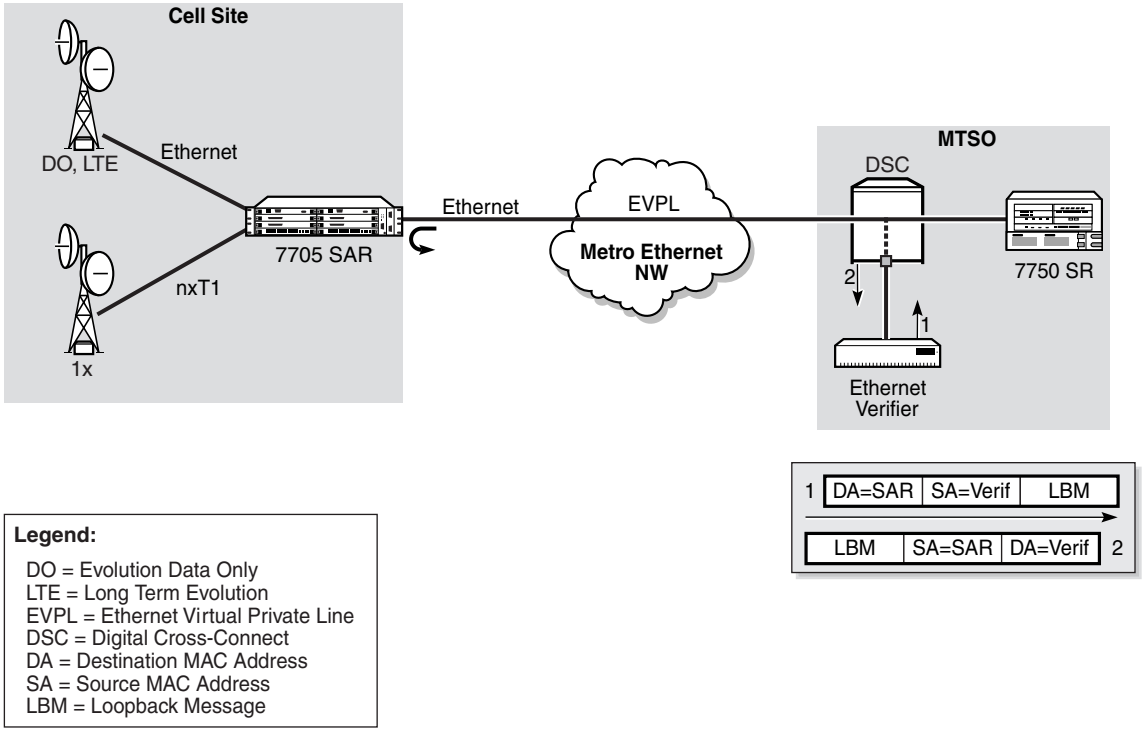

Figure 9 shows an Ethernet verifier at the MTSO that is directly connected to the transport network (in front of the 7750 SR). Thus, the Ethernet OAM frames are not label-encapsulated. Given that Ethernet verifiers do not support label operations and the transport provider mandates that OAM tests be run between the two hand-off Ethernet ports, the verifier cannot be relocated behind the 7750 SR node at the MTSO. Therefore, CFM loopback frames received are not MPLS-encapsulated, but are simple Ethernet frames where the `type` is set to CFM (dot1ag or Y.1731).

CFM Loopback Mechanics

The following list contains important facts to consider when working with CFM loopbacks:

- CFM loopbacks can be enabled on a per-port basis, and:
  - → the port can be in access or network mode
  - → once enabled on a port, all received LBM frames are processed, regardless of the VLAN and the service that the VLAN or SAP is bound to
  - → there is no associated MEP creation involved with this feature; therefore, no domain, association, or similar checks are performed on the received frame
  - → upon finding a destination address MAC match, the LBM frame is sent to the CFM process
- received LBM frames undergo no queuing or scheduling in the ingress direction
- at egress, loopback reply (LBR) frames are stored in their own queue; that is, a separate new queue is added exclusively for LBR frames
- users can configure the way a response frame is treated among other user traffic stored in network queues; the configuration options are high-priority or low-priority
- for network egress, where profiled scheduling is enabled, the following conditions apply:
  - → **high-priority**: either cir = port_speed, which applies to all frames that are scheduled via an in-profile scheduler; or round-robin (RR) for all other (network egress queue) frames that are in-profile
  - → **low-priority**: either cir = 0, pir = port_speed, which applies to all frames that are scheduled as out-of-profile, or RR for all other frames that are out-of-profile
- for network egress or access egress, where 4-priority scheduling is enabled:
  - → **high-priority**: either cir = port_speed, which applies to all frames that are scheduled via an expedited in-profile scheduler, or RR for all other (network egress queue) frames that reside in expedited queues and are in an in-profile state
  - → **low-priority**: either cir = 0, pir = port_speed, which applies to all frames that are scheduled via a best effort out-of-profile scheduler, or RR for all other frames that reside in best-effort queues and are in an out-of-profile state
- the above queue parameters and scheduler mappings are all preconfigured and cannot be altered. The desired QoS treatment is selected by enabling the CFM loopback and specifying high priority or low priority.

OAM Propagation to Attachment Circuits

Typically, T1/E1 equipment at a site relies on the physical availability of the T1/E1 ports to determine the uplink capacity (that is, for ATM IMA or MLPPP groups). When a failure in the access link between the 7705 SAR and the associated T1/E1 equipment is detected, notification of the failure is propagated by the PW status signaling using one of two methods — label withdrawal or TLV (see LDP Status Signaling). In addition, the PW failure must also be propagated to the devices attached to the T1/E1 equipment. The propagation method depends on the type of port used by the access circuit (ATM, T1/E1 TDM, or Ethernet) and is described below.

ATM Ports

Propagation of ATM PW failures to the ATM port is achieved through the generation of AIS and RDI alarms.

T1/E1 TDM Ports

If a port on a 16-port T1/E1 ASAP Adapter card, 32-port T1/E1 ASAP Adapter card, or 2-port OC3/STM1 Channelized Adapter card is configured for CESoPSN VLL service, failure of the VLL forces a failure of the associated DS0s (timeslots). Since there can be $n \times$ DS0s bound to a CESoPSN VLL service as the attachment circuit, an alarm is propagated to the bound DS0s only. In order to emulate the failure, an “all 1s” or an “all 0s” signal is sent through the DS0s. The bit pattern can be configured to be either all 1s or all 0s. This is sometimes called “trunk conditioning”.

Ethernet Ports

For an Ethernet port-based Ethernet VLL, failure of the VLL forces a failure of the local Ethernet port. That is, the local attachment port is taken out of service at the physical layer and the Tx is turned off on the associated Ethernet port.

LDP Status Signaling

The failure of a local circuit needs to be propagated to the far-end PE, which then propagates the failure to its attached circuits. The 7705 SAR can propagate failures over the PW using one of the following methods:

- LDP status via label withdrawal
- LDP status via TLV

LDP Status via Label Withdrawal

Label withdrawal is negotiated during the PW status negotiation phase and needs to be supported by both the near-end and the far-end points. If the far end does not support label withdrawal, the 7705 SAR still withdraws the label in case the local attachment circuit is removed or shut down.

Label withdrawal occurs only when the attachment circuit is administratively shut down or deleted. If there is a failure of the attached circuit, the label withdrawal message is not generated.

When the local circuit is re-enabled after shutdown, the VLL must be re-established, which causes some delays and signaling overhead.

LDP Status via TLV

Signaling PW status via TLV is supported as per RFC 4447. Signaling PW status via TLV is advertised during the PW capabilities negotiation phase. It is more efficient and is preferred over the label withdrawal method.

For cell mode ATM PWs, when an AIS message is received from the local attachment circuit, the AIS message is propagated to the far-end PE unaltered and PW status TLV is not initiated.

Service Assurance Agent (SAA) Overview

Broadband service delivery technologies have enabled the introduction of broadband service termination applications such as Voice over IP (VoIP), TV service delivery, and video and high-speed Internet services. These new applications force carriers to produce services where the health and quality of Service Level Agreement (SLA) commitments are verifiable, both to the customer and internally (within the carrier).

SAA is a feature that monitors network operations using statistics for parameters such as latency, jitter, response time, and packet loss. The information can be used to troubleshoot network problems and help in problem prevention and network topology planning. The 7705 SAR also supports the following SAA Ethernet CFM tests: loopback, linktrace, and two-way delay measurement.

The results are saved in SNMP tables that 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 CFM tests can be saved to accounting files that can be accessed by the network management system.

SAA Application

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

For SAA ICMP ping, one-way timestamping can be enabled at the system level for all outbound SAA ICMP ping packets.

Traceroute Implementation

For various applications, such as LSP traceroute, packets must pass through the network processor while on their way to the control CPU. When the packets exit the control CPU in the egress direction, the network processor inserts a timestamp inside each packet. Only packets processed by the control CPU will receive a timestamp.

When interpreting these timestamps, care must be taken because 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 which hop is being measured.

SAA Jitter

Mobile operators require millisecond-level granularity when it comes to delay and jitter measurements. This is especially true for synchronization-over-packet based applications.

Two-way jitter tests measure the jitter in each direction separately. The 7705 SAR provides two-way jitter tests with millisecond granularity for all network deployment applications.

SAA Ethernet CFM Test Support

CFM loopback, linktrace, and two-way delay measurement (Y.1731 ETH-DMM) tests can be initiated using SAA. Additional timestamping is required for 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 that any time measurements resulting from the loopback and linktrace tests include the packet processing time used by the remote node. Since Y.1731 ETH-DMM uses a four timestamp 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-DDM tests support send-count, interval, timeout, and FC. The summary of the test results are stored in an SAA accounting file.

Writing SAA Ethernet CFM Test Results to Accounting Files

When each SAA CFM test is completed, the 7705 SAR collects the results in an accounting file that can be accessed by the network management system. In order to write the SAA test results to an accounting file in a compressed XML format, the results must be collected and entered in the appropriate MIB table, and a record must be generated in the appropriate accounting file.

Refer to the 7705 SAR OS System Management Guide, “Configuring an Accounting Policy” section, for information about creating accounting files and writing to them.

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

Configuring SAA Test Parameters

Use the following CLI syntax to create an SAA test and set test parameters:

CLI Syntax:
```
config# saa
config>saa# test ping
config>saa>test$ type
config>saa>test>type$ icmp-ping 10.10.221.131 count 50 fc
"nc" profile out
config>saa>test>type$ exit
config>saa>test# no shutdown
config>saa>test# exit
config>saa# exit
```

The following example displays the SAA test configuration output:

```
A:ALU-48>config>saa
-----------------------------------------------
        test "ping"
            type
                icmp-ping 10.10.221.131 count 50 fc "nc" profile out
            exit
            no shutdown
        exit
-----------------------------------------------
```

The following example displays the result after running the test:

```
A:ALU-48>config>saa# show saa ping
===============================================================================
SAA Test Information
===============================================================================
Test name                    : ping
Owner name                   : TiMOS CLI
Description                  : N/A
Accounting policy            : None
Administrative status        : Enabled
Test type                    : icmp-ping 10.10.221.131 count 50 fc "nc"
                               profile out
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: 50
Number of requests that failed to be sent out: 0
Number of responses that were received: 50
Number of requests that did not receive any response: 0
Total number of failures: 0, Percentage: 0
 (in ms)            Min          Max      Average       Jitter
Outbound  :        -9.61        -8.75        -9.18        0.016
Inbound   :         9.53         12.0         10.2        0.412
Roundtrip :        0.674         2.59         1.02        0.406

Per test packet:
  Sequence     Outbound      Inbound    RoundTrip Result
         1        -8.75         9.53        0.784 Response Received
         2        -8.76         9.54        0.779 Response Received
         3        -8.78         9.59        0.805 Response Received
         4        -8.79         11.3         2.46 Response Received
         5        -8.82         9.61        0.786 Response Received
         6        -8.83         9.59        0.760 Response Received
         7        -8.86         9.65        0.795 Response Received
         8        -8.86         9.63        0.767 Response Received
         9        -8.89         9.68        0.797 Response Received
        10        -8.90         9.68        0.775 Response Received
        11        -8.93         9.73        0.805 Response Received
        12        -8.93         10.4         1.44 Response Received
        13        -8.97         9.75        0.788 Response Received
        14        -8.98         11.2         2.23 Response Received
        15        -9.00         9.80        0.801 Response Received
        16        -9.01         9.79        0.787 Response Received
        17        -9.03         9.82        0.794 Response Received
        18        -9.04         10.9         1.89 Response Received
        19        -9.06         9.87        0.801 Response Received
        20        -9.08         9.85        0.770 Response Received
        21        -9.10         9.90        0.804 Response Received
        22        -9.11         9.90        0.782 Response Received
        23        -9.14         9.97        0.828 Response Received
        24        -9.15         9.93        0.780 Response Received
        25        -9.17         9.99        0.813 Response Received
        26        -9.18         9.97        0.786 Response Received
        27        -9.21         10.5         1.28 Response Received
        28        -9.22         11.0         1.79 Response Received
        29        -9.25         10.1        0.807 Response Received
        30        -9.26         10.0        0.767 Response Received
        31        -9.28         10.1        0.804 Response Received
        32        -9.29         9.96        0.676 Response Received
        33        -9.31         10.0        0.719 Response Received
```

```
        34        -9.32         10.1        0.785 Response Received
        35        -9.35         10.2        0.808 Response Received
        36        -9.36         10.1        0.782 Response Received
        37        -9.39         11.3         1.87 Response Received
        38        -9.40         12.0         2.59 Response Received
        39        -9.43         10.2        0.792 Response Received
        40        -9.43         10.2        0.771 Response Received
        41        -9.46         10.3        0.815 Response Received
        42        -9.46         10.1        0.674 Response Received
        43        -9.49         12.0         2.46 Response Received
        44        -9.50         10.3        0.782 Response Received
        45        -9.53         10.3        0.810 Response Received
        46        -9.54         10.3        0.780 Response Received
        47        -9.57         10.3        0.768 Response Received
        48        -9.58         10.3        0.769 Response Received
        49        -9.60         10.4        0.797 Response Received
        50        -9.61         11.2         1.60 Response Received

===============================================================================
*A:ALU-48#
```

OAM and SAA Command Reference

Command Hierarchies

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*] [**interval** *interval*] [{**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** *ip-address*] [**tos** *type-of-service*] [**router** *router-instance*]

OAM Commands

ATM Diagnostics

global

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

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**]

SDP Diagnostics

global

- — **oam**
  - — **sdp-mtu** *orig-sdp-id* **size-inc** *start-octets end-octets* [**step** *step-size*] [**timeout** *timeout*] [**interval** *interval*]
  - — **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Service Diagnostics

global
— **oam**
— **svc-ping** *ip-address* **service** *service-id* [**local-sdp**] [**remote-sdp**]
— **vprn-ping** *service-id* **source** *ip-address* **destination** *ip-address* [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *size*] [**ttl** *vc-label-ttl*] [**count** *send-count*] [**return-control**] [**timeout** *timeout*] [**interval** *interval*]
— **vprn-trace** *service-id* **source** *ip-address* **destination** *ip-address* [**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*]

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*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**ttl** *vc-label-ttl*]
— **vccv-trace** *sdp-id:vc-id* [**size** *octets*] [**min-ttl** *min-vc-label-ttl*] [**max-ttl** *max-vc-label-ttl*] [**max-fail** *no-response-count*] [**probe-count** *probe-count*] [**reply-mode** {**ip-routed** | **control-channel**}] [**timeout** *timeout-value*] [**interval** *interval-value*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**detail**]

VPLS Diagnostics

global
— **oam**
— **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*] [**count** *send-count*] [**send-control**] [**return-control**] [**timeout** *timeout*] [**interval** *interval*]
— **mac-ping** **service** *service-id* **destination** *dst-ieee-address* [**source** *src-ieee-address*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**ttl** *vc-label-ttl*] [**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}
            — remote-loopback {start | stop}

config
    — [no] port {port-id}
        — ethernet
            — efm-oam
                — [no] accept-remote-loopback
                — hold-time time-value
                — no hold-time
                — mode {active | passive}
                — [no] shutdown
                — [no] transmit-interval interval [multiplier multiplier]
                — [no] tunneling
```

ETH-CFM Commands

global
- **oam**
  - **eth-cfm**
    - **eth-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*] [**data-length** *data-length*]
    - **linktrace** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**ttl** *ttl-value*]
    - **loopback** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**send-count** *send-count*] [**size** *data-size*] [**priority** *priority*]
    - **one-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]
    - **single-ended-loss-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*] [**interval** {**100ms** | **1s**}] [**send-count** *send-count*]
    - **two-way-delay-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*]

config
- **eth-cfm**
  - **domain** *md-index* [**format** {**dns** | **mac** | **none** | **string**}] **name** *md-name* **level** *level*
  - **domain** *md-index*
  - **no domain** *md-index*
    - **association** *ma-index* [**format** {**icc-based** | **integer** | **string** | **vid** | **vpn-id**}] **name** *ma-name*
    - **association** *ma-index*
    - **no association** *ma-index*
      - [**no**] **bridge-identifier** *bridge-id*
        - **vlan** *vlan-id*
        - **no vlan**
      - **ccm-interval** {**10ms** | **100ms** | **1** | **10** | **60** | **600**}
      - **no ccm-interval**
    - [**no**] **remote-mepid** *mep-id*

config
- [**no**] **port** {*port-id*}
  - **ethernet**
    - **cfm-loopback priority** {**low** | **high**}
    - **no cfm-loopback**

config
- **service**
  - [**no**] **epipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*]
    - **sap** *sap-id* [**create**]
      - **eth-cfm**
        - **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up** | **down**}]
        - **no mep** *mep-id* **domain** *md-index* **association** *ma-index*
          - [**no**] **ais-enable**
            - **client-meg-level** [*level* [*level ...*]]
            - [**no**] **client-meg-level**
            - **interval** [**1** | **60**]
            - [**no**] **interval**
            - **priority** *priority-value*

```
                                        — [no] priority
                                    — [no] ccm-enable
                                    — ccm-ltm-priority priority
                                    — no ccm-ltm-priority
                                    — [no] dual-ended-loss-test-enable
                                        — alarm-threshold percentage
                                        — no alarm-threshold
                                        — alarm-clear-threshold percentage
                                        — no alarm-clear-threshold
                                    — [no] eth-test-enable
                                        — bit-error-threshold bit-errors
                                        — [no] test-pattern {all-zeros | all-ones} [crc-
                                          enable]
                                    — low-priority-defect {allDef | macRemErrXcon |
                                      remErrXcon | errXcon | xcon | noXcon}
                                    — one-way-delay-threshold seconds
                                    — [no] shutdown
                            — spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}] [create] [no-endpoint]
                            — spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}] [create] endpoint endpoint-name
                                — eth-cfm
                                    — mep mep-id domain md-index association ma-index
                                      [direction {up | down}]
                                    — no mep mep-id domain md-index association ma-index
                                        — [no] ccm-enable
                                        — ccm-ltm-priority priority
                                        — [no] ccm-ltm-priority
                                        — low-priority-defect {allDef | macRemErrXcon |
                                          remErrXcon | errXcon | xcon | noXcon}
                                        — [no] shutdown
```

Configure SAA Commands

```
config
    — saa
        — [no] test test-name [owner test-owner]
            — description description-string
            — no description
            — jitter-event rising-threshold threshold [falling-threshold threshold] [direction]
            — no jitter-event
            — latency-event rising-threshold threshold [falling-threshold threshold]
              [direction]
            — no latency-event
            — loss-event rising-threshold threshold [falling-threshold threshold] [direction]
            — no loss-event
            — [no] shutdown
            — [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] [count send-count] [send-control] [return-control] [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**}]] [**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**}]] [**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* [**profile** {**in** | **out**}]] [**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*] [**interval** *seconds*] [{**next-hop** *ip-address*} | {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance*] [**timeout** *timeout*] [**fc** *fc-name* [**profile** {**in** | **out**}]]
- **icmp-trace** [*ip-address* | *dns-name*] [**ttl** *time-to-live*] [**wait** *milli-seconds*] [**source** *ip-address*] [**tos** *type-of-service*] [**router** *router-instance*]
- **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*]]
- **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*]]
- **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*] [**count** *send-count*] [**send-control**] [**return-control**] [**interval** *interval*] [**timeout** *timeout*]
- **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*]
- **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]
- **vccv-ping** *sdp-id:vc-id* [**src-ip-address** *ip-addr* **dst-ip-address** *ip-addr* **pw-id** *pw-id*] [**reply-mode** {**ip-routed** | **control-channel**}] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**ttl** *vc-label-ttl*]
- **vccv-trace** *sdp-id:vc-id* [**size** *octets*] [**min-ttl** *min-vc-label-ttl*] [**max-ttl** *max-vc-label-ttl*] [**max-fail** *no-response-count*] [**probe-count** *probe-count*] [**reply-mode** {**ip-routed** | **control-channel**}] [**timeout** *timeout-value*] [**interval** *interval-value*] [**fc** *fc-name* [**profile** {**in** |**out**}]] [**detail**]
- **vprn-ping** *service-id* **source** *ip-address* **destination** *ip-address* [**fc** *fc-name* [**profile** {**in** | **out**]] [**size** *size*] [**ttl** *vc-label-ttl*] [**count** *send-count*] [**return-control**] [**timeout** *timeout*] [**interval** *interval*]
- **vprn-trace** *service-id* **source** *ip-address* **destination** *ip-address* [**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*]

```
config
    — system
        — [no] enable-icmp-vse
```

SAA Diagnostics

```
global
    — oam
        — saa test-name [owner test-owner] {start | stop}
```

Show Commands

```
show
    — eth-cfm
        — association [ma-index] [detail]
        — cfm-stack-table
        — cfm-stack-table port [port-id [vlan vlan-id]] [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]
        — 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-meps}
        — 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 single-ended-loss-test [remote-peer mac-address]
        — mep mep-id domain md-index association ma-index dual-ended-loss-test [remote-peer mac-address]
    — saa [test-name [owner test-owner]]
    — ldp-treetrace [prefix ip-prefix/mask] [detail]
```

Clear Commands

```
clear
    — saa [test-name [owner test-owner]]
    — eth-cfm
        — dual-ended-loss-test mep mep-id domain md-index association ma-index
```

Debug Commands

```
debug
    — [no] oam
            — lsp-ping-trace [tx | rx | both] [raw | detail]
            — no lsp-ping-trace
```

Command Descriptions

OAM and SAA Commands

Operational Commands

ping

Syntax **ping** [*ip-address* | *dns-name*] [**rapid** | **detail**] [**ttl** *time-to-live*] [**tos** *type-of-service*] [**size** *bytes*] [**pattern** *pattern*] [**source** *ip-address*] [**interval** *interval*] [{**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* — identifies the far-end IP address to which to send the **ping** request message

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 to FFFF]H d: [0 to 255]D

dns-name — identifies the DNS name of the far-end device to which to send the **ping** request message, expressed as a character string

rapid — specifies that packets will be generated as fast as possible instead of the default 1 per second

detail — displays detailed information

time-to-live — specifies the TTL value for the MPLS label, expressed as a decimal integer

Values 1 to 128

type-of-service — specifies the service type

Values 0 to 255

bytes **—** specifies the request packet size in bytes, expressed as a decimal integer

Values 0 to 16384

pattern — specifies the pattern that will be used to fill the date portion in a ping packet. If no pattern is specified, position information will be filled instead.

Values 0 to 65535

source *ip-address* **—** specifies the IP address to be used

Values

ipv4-address	a.b.c.d (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 to FFFF]H d: [0 to 255]D

interval — defines the minimum amount of time, expressed as a decimal integer, that must expire before the next message request is sent.

This parameter is used to override the default request message send interval. If the **interval** is set to 1 s, and the **timeout** value is set to 10 s, then the maximum time between message requests is 10 s and the minimum is 1 s. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Values 1 to 10

Default 1

next-hop *ip-address* **—** displays only the 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 to FFFF]H d: [0 to 255]D

interface-name — specifies the name of an IP 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

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

Values 1 to 100000

Default 5

do-not-fragment — sets the DF (Do not fragment) bit in the ICMP ping packet (does not apply to ICMPv6)

router-instance — specifies the router name or service ID

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

Default Base

timeout — specifies 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.

This value is used to override the default timeout value.

Values 1 to 10

Default 5

shutdown

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

Context config>saa>test
config>port>ethernet>efm-oam
config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description The **shutdown** command administratively disables a test. A **shutdown** can only be performed if a test is not executing at the time the command is entered.

When a test is created, it remains in shutdown mode until a **no shutdown** command is executed.

In order to modify an existing test, it must first be shut down.

When used with the **ethernet>efm-oam** command, **shutdown** enables tunneling on the port (see tunneling), and **no shutdown** enables Ethernet EFM OAM 802.3ah.

The **no** form of this command sets the state of the test to operational.

Default **shutdown**

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 <GLOBAL>

Description This command determines the route to a destination address.

Parameters *ip-address* — specifies the far-end IP address to which to send the traceroute request message

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 to FFFF]H d: [0 to 255]D

dns-name — specifies the DNS name of the far-end device to which to send the traceroute request message, expressed as a character string

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

Values 1 to 255

milli-seconds — specifies the time in milliseconds to wait for a response to a probe, expressed as a decimal integer

Values 10 to 60000

Default 5000

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 of the responding hosts are performed

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

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 to FFFF]H d: [0 to 255]D

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

Values 0 to 255

router-instance — specifies a router name or service ID

Values	router-name	Base, management
	service-id	1 to 2147483647

Default Base

Output **Sample Destination Address Route**

```
*A:ALU-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:ALU-1#
```

ATM Diagnostics

atm-ping

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

Context oam

Description This command tests ATM path connectivity 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*	
bundle-id	bundle-*type-slot/mda.bundle-num*	
	bundle	keyword
	type	ima
	bundle-num	1 to 10
vpi	0 to 4095 (NNI)	
	0 to 255 (UNI)	
vci	1, 2, 5 to 65535	

end-to-end | segment — specifies whether the ATM OAM loopback cell is destined for the first segment point in the line direction or the PVCC's connection endpoint

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 0x11 will be used

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 time out or receive a reply before the next message request is sent. The message interval value must be expired before the next message request is sent.

Values 1 to 100

Default 1

timeout — specifies 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.

This value is used to override the default timeout value.

Values 1 to 10

Default 5

interval — specifies the minimum amount of time that must expire before the next message request is sent

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

This parameter is used to override the default request message send interval.

Values 1 to 10

Default 1

Service Diagnostics

sdp-mtu

Syntax **sdp-mtu** *orig-sdp-id* **size-inc** *start-octets end-octets* [**step** *step-size*] [**timeout** *timeout*] [**interval** *interval*]

Context oam

Description This command performs MTU path tests on an SDP to determine the largest **path-mtu** supported on an SDP. The **size-inc** parameter can be used to easily determine the **path-mtu** of a given SDP-ID. The forwarding class is assumed to be Best-Effort Out-of-Profile. The message reply is returned with IP encapsulation from the far-end 7705 SAR. OAM request messages sent within an IP SDP must have the "DF" IP header bit set to 1 to prevent message fragmentation.

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 a response timeout occurs, the maximum size message replied to indicates the largest size OAM request message that received a valid reply.

To terminate an **sdp-mtu** in progress, use the CLI break sequence <Ctrl-C>.

Parameters *orig-sdp-id* — specifies the SDP-ID to be used by **sdp-mtu,** 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 SDP tunnel encapsulation used to reach the far end — GRE, IP, or MPLS. If *orig-sdp-id* is invalid or administratively down or unavailable for some reason, the SDP echo request message is not sent and an appropriate error message is displayed (once the **interval** timer expires, **sdp-mtu** will attempt to send the next request if required).

Values 1 to 17407

start-octets end-octets **—** indicates that an incremental path MTU test will be performed by sending a series of message requests with increasing MTU sizes

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

Values 40 to 9198

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

Values 40 to 9198

step-size — specifies 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.

Values 1 to 512

Default 32

timeout — specifies 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.

This value is used to override the default **timeout** value.

Values 1 to 10

Default 5

interval — defines the minimum amount of time that must expire before the next message request is sent.

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

This parameter is used to override the default request message send interval.

Values 1 to 10

Default 1

Output **Sample SDP MTU Path Test Output**

```
*A:router 1> sdp-mtu 6 size-inc 512 3072 step 256
   Size        Sent      Response
 -------       ----      --------------------------
   512         .         Success
   768         .         Success
  1024         .         Success
  1280         .         Success
  1536         .         Success
  1792         .         Success
  2048         .         Success
  2304         …         Request Timeout
  2560         …         Request Timeout
  2816         …         Request Timeout
  3072         …         Request Timeout
Maximum Response Size:   2048
```

svc-ping

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

Context oam

Description This command tests a service ID for correct and consistent provisioning between two service endpoints. The command accepts a far-end IP address and a Service-ID for local and remote service testing. The following information can be determined from **svc-ping**:

- local and remote service existence
- local and remote service state
- local and remote service type correlation
- local and remote customer association
- local and remote service-to-SDP bindings and state
- local and remote ingress and egress service label association

Unlike **sdp-ping**, only a single message will be sent per command; no count or interval parameter is supported and round-trip time is not calculated. A timeout value of 10 s 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 an **svc-ping** in progress, use the CLI break sequence <Ctrl-C>.

Upon request timeout, message response, request termination, or request error, the local and remote information described in Table 4 will be displayed. Local and remote information is dependent upon service existence and reception of reply.

Table 4: SVC Ping Report Field

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

Table 4: SVC Ping Report Field (Continued)

Field	Description	Values
Service-ID	The Service-ID 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, 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
Remote Service Type	The remote type of service being tested. If *service-id* does not exist remotely, N/A is displayed.	Epipe, 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

Table 4: SVC Ping Report Field (Continued)

Field	Description	Values
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 a 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. The **sdp-ping** command should also fail.	resp-ip-addr N/A
Responders Expected Far-end Address	The expected source of the originator's SDP-ID from the perspective of the remote 7705 SAR 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

Table 4: SVC Ping Report Field (Continued)

Field	Description	Values
Originating SDP-ID Path Used	Indicates whether the originating 7705 SAR used the originating SDP-ID to send the **svc-ping** request. If a valid originating SDP-ID is found, is operational and has a valid egress service label, the originating 7705 SAR should use the SDP-ID as the requesting path if **sdp-path** has been defined. If the originating 7705 SAR uses the originating SDP-ID as the request path, Yes is displayed. If the originating 7705 SAR 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 shut down, 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-Down 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-Down 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 7705 SAR 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 7705 SAR that is bound to the *service-id*, Non-Existent is displayed.	resp-sdp-id Non-Existent

Table 4: SVC Ping Report Field (Continued)

Field	Description	Values
Responding SDP-ID Path Used	Indicates whether the responding 7705 SAR 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, is operational and has a valid egress service label, the far-end 7705 SAR 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-Down 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

Table 4: SVC Ping Report Field (Continued)

Field	Description	Values
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
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 7705 SAR considers the displayed egress service label operational, Up is displayed. If the originating 7705 SAR 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 7705 SAR for this *service-id* to the originating 7705 SAR. If *service-id* does not exist in the remote 7705 SAR, 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 considers its egress service label operational, Up is displayed. If the responding 7705 SAR 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 7705 SAR. 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

Table 4: SVC Ping Report Field (Continued)

Field	Description	Values
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 originator's ingress service label has not been assigned, N/A is displayed.	Manual Signaled N/A
Expected Ingress Service Label State	The originator's ingress service label state. If the originating 7705 SAR considers its ingress service label operational, Up is displayed. If the originating 7705 SAR 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 7705 SAR. This is the service label that the far end is expecting to receive for *service-id* when sending to the originating 7705 SAR. If *service-id* does not exist in the remote 7705 SAR, N/A is displayed. If *service-id* exists, but an ingress service label has not been assigned in the remote 7705 SAR, 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 7705 SAR. If the ingress service label is manually defined on the remote 7705 SAR, Manual is displayed. If the ingress service label is dynamically signaled on the remote 7705 SAR, Signaled is displayed. If the *service-id* does not exist on the remote 7705 SAR, N/A is displayed.	Manual Signaled N/A
Responders Ingress Service Label State	The assigned ingress service label state on the remote 7705 SAR. If the remote 7705 SAR considers its ingress service label operational, Up is displayed. If the remote 7705 SAR considers its ingress service label inoperative, Down is displayed. If the *service-id* does not exist on the remote 7705 SAR or the ingress service label has not been assigned on the remote 7705 SAR, N/A is displayed.	Up Down N/A

Parameters *ip-address* — specifies the far-end IP address to which to send the **svc-ping** request message in dotted-decimal notation

service-id — identifies the service being tested. The Service ID need not exist on the local 7705 SAR to receive a reply message.

This is a mandatory parameter.

Values 1 to 2147483647

local-sdp — specifies that 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 SDP tunnel encapsulation used to reach the far end — GRE, IP, or MPLS. On originator egress, the service-ID must have an associated VC-Label to reach the far-end address of the SDP-ID and the SDP-ID must be operational for the message to be sent.

If **local-sdp** is not specified, the **svc-ping** request message is sent with GRE encapsulation with the OAM label.

Table 5 indicates whether a message is sent and how the message is encapsulated based on the state of the service ID.

Table 5: Local SDP Message Results

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 that the **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 SDP tunnel encapsulation used to reply to the originator — GRE, IP, or MPLS. On responder egress, the service-ID must have an associated VC-Label to reach the originator address of the SDP-ID and the SDP-ID must be operational for the message to be sent. If **remote-sdp** is not specified, the **svc-ping** request message is sent with GRE encapsulation with the OAM label.

Table 6 indicates how the message response is encapsulated based on the state of the remote Service ID.

Table 6: Remote SDP Message Results

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
Service-ID: 101

Err Info                Local           Remote
------------------------------------------------------
    Type:               CPIPE           CPIPE
    Admin State:        Up              Up
    Oper State:         Up              Up
    Service-MTU:        1000            1000
    Customer ID:        1001            1001

==> IP Interface State: Down
    Actual IP Addr:     10.10.10.11     10.10.10.10
    Expected Peer IP:   10.10.10.10     10.10.10.11

==> SDP Path Used:      Yes             Yes
    SDP-ID:             123             325
    Admin State:        Up              Up
    Operative State:    Up              Up
    Binding Admin State:Up              Up
    Binding Oper State: Up              Up
    Binding VC ID:      101             101
    Binding Type:       Spoke           Spoke
    Binding Vc-type:    CesoPsn         CesoPsn
    Binding Vlan-vc-tag:0               0
```

```
==> Egress Label:       131066          131064
    Ingress Label:      131064          131066
    Egress Label Type:  Signaled        Signaled
    Ingress Label Type: Signaled        Signaled

Request Result: Sent - Reply Received
```

EFM Commands

efm

Syntax **efm** *port-id*

Context oam

Description This command enables Ethernet in the First Mile (EFM) OAM loopbacks on the specified port. The EFM OAM remote loopback OAMPDU will be sent to the peering device to trigger a remote loopback.

Parameters *port-id* — specifies 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 EFM OAM loopback tests on the specified port. The EFM OAM remote loopback OAMPDU will be sent to the peering device to trigger a remote loopback.

ethernet

Syntax **ethernet**

Context config>port

Description This command enables access to the context to configure Ethernet port attributes on an 8-port Ethernet Adapter card.

efm-oam

Syntax **efm-oam**

Context config>port>ethernet

Description This command configures EFM OAM attributes.

accept-remote-loopback

Syntax [**no**] **accept-remote-loopback**

Context config>port>ethernet>efm-oam

Description This command enables reactions to loopback control OAMPDUs from peers.

The **no** form of this command disables reactions to loopback control OAMPDUs.

hold-time

Syntax **hold-time** *time-value*
no hold-time

Context config>port>ethernet>efm-oam

Description This command sets the amount of time that EFM-OAM will wait before going from a non-operational state to an operational state.

If EFM-OAM goes from an operational state to a non-operational state (other than link-fault), it enters the hold-time period. During this time, EFM-OAM continues to negotiate with the peer if possible, but will not transition to the "up" state until the hold time has expired.

If EFM-OAM goes down due to a lower-level fault (for example, the port goes down and EFM-OAM enters the link-fault state), the hold timer is not triggered. When the lower-level fault is cleared, EFM-OAM immediately starts running on the port and transitions to the operational state as soon as possible.

If EFM-OAM goes down because the user administratively disables the protocol, EFM-OAM immediately transitions to the disabled state. When the user re-enables EFM-OAM, the protocol enters the hold time period and EFM-OAM is not operational until the hold time expires.

A hold-time value of 0 indicates that EFM-OAM returns to the operational state without delay.

The hold time affects only the transition from a non-operational state to an operational state; it does not apply to a transition from an operational state to a non-operational state.

Parameters *time-value* — the number of seconds that EFM-OAM will wait before returning to an operational state from a non-operational state

Values 0 to 50

Default 0

mode

Syntax **mode** {**active** | **passive**}

Context config>port>ethernet>efm-oam

Description This command configures the mode of OAM operation for this Ethernet port.

Active mode causes the port to initiate the negotiation process and continually send out EFM OAM information PDUs. **Passive** mode waits for the peer to initiate the negotiation process. A passive mode port cannot initiate monitoring activities (such as loopback) with the peer.

Default **active**

transmit-interval

Syntax [**no**] **transmit-interval** *interval* [**multiplier** *multiplier*]

Context config>port>ethernet>efm-oam

Description This command configures the transmit interval of OAMPDUs.

Parameters *interval* — specifies the transmit interval

Values 1 to 600 (in 100 ms)

multiplier — specifies the multiplier for the transmit interval to set the local link down timer

Values 2 to 5

tunneling

Syntax [**no**] **tunneling**

Context config>port>ethernet>efm-oam

Description This command enables EFM OAMPDU tunneling. OAMPDU tunneling is required when a loopback is initiated from a router end and must be transported over the existing network infrastructure to the other end. Enabling tunneling will allow the PDUs to be mapped to Epipes so that the OAM frames can be tunneled over MPLS to the far end.

To enable Ethernet EFM OAM 802.3ah on the port, use the **efm-oam>no shutdown** command.

The **no** form of the command disables tunneling.

ETH-CFM Commands

eth-test

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

Context oam eth-cfm

Description This command specifies to initiate an Ethernet (signal) test.

Parameters *mac-address* — specifies a unicast MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the MD index

Values 1 to 4294967295

ma-index — specifies the MA index

Values 1 to 4294967295

priority — specifies the value used for priority mapping

Values 0 to 7

Default the CCM and LTM priority of the MEP

data-length — specifies the packet size in bytes, expressed as a decimal integer, used for the ETH-CFM test

Values 64 to 1500

Default 64

linktrace

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

Context oam>eth-cfm

Description This command specifies to initiate a linktrace test.

Parameters *mac-address* — specifies a unicast destination MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the MD index

Values 1 to 4294967295

ma-index — specifies the MA index

Values 1 to 4294967295

ttl-value — specifies the TTL for a returned linktrace

Values 0 to 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

Description This command specifies to initiate a loopback test.

Parameters *mac-address* — specifies a unicast MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the MD index

Values 1 to 4294967295

ma-index — specifies the MA index

Values 1 to 4294967295

send-count — specifies the number of messages to send, expressed as a decimal integer. Dot1ag loopback messages are sent back-to-back, with no delay between the transmissions.

Values 1 to 5

Default 1

data-size — specifies the packet size in bytes, expressed as a decimal integer

Values 0 to 1500

Default 0

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

Values 0 to 7

Default the CCM and LTM priority of the MEP

one-way-delay-test

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

Context oam>eth-cfm

Description This command specifies to initiate an ETH-CFM one-way delay test.

Parameters *mac-address* — specifies a unicast MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the MD index

Values 1 to 4294967295

ma-index — specifies the MA index

Values 1 to 4294967295

priority — specifies the value used for priority mapping

Values 0 to 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 specifies to initiate an ETH-CFM two-way delay test.

The *priority* is selected according to the mappings in Table 7.

Table 7: Y.1731 Priority-to-FC Mapping

Priority	FC-ID	FC Name
0	0	BE
1	1	L2
2	2	AF
3	3	L1
4	4	H2

Table 7: Y.1731 Priority-to-FC Mapping (Continued)

Priority	FC-ID	FC Name
5	5	EF
6	6	H1
7	7	NC

Parameters *mac-address* — specifies a unicast MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the MD index

Values 1 to 4294967295

ma-index — specifies the MA index

Values 1 to 4294967295

priority — specifies the priority mapping value that specifies the FC for OAM traffic, according to Table 7

Values 0 to 7

Default The CCM and LTM priority of the MEP

single-ended-loss-test

Syntax **single-ended-loss-test** *mac-address* **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**priority** *priority*] [**interval** {**100ms** | **1s**}] [**send-count** *send-count*]

Context oam>eth-cfm

Description This command specifies to initiate a loss measurement test between the specified *mac-address* router and the specified *mep-id* MEP.

Single-ended and dual-ended loss tests are mutually exclusive tests. Single-ended loss tests can be run when dual-ended loss tests are disabled (under the **spoke-sdp>eth-cfm>mep** context).

Parameters *mac-address* — specifies a unicast MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the index of the MD to which the MEP is associated, or 0, if none

Values 1 to 4294967295

ma-index — specifies the index to which the MEP is associated, or 0, if none

Values 1 to 4294967295

send-count — specifies the number of LMM messages to send, expressed as a decimal integer

Values 2 to 5

Default 2

interval {**100ms | 1s**} — specifies the interval between groups of consecutive LMM packets (for example, if *send-count* is 5 and *interval* is 1s, then 5 LMM packets are sent at 1-s intervals)

Values 100ms | 1s

Default 1s

priority — specifies the priority mapping value that specifies the FC for OAM traffic, according to Table 7

Values 0 to 7

Default the CCM and LTM priority of the MEP

eth-cfm

Syntax **eth-cfm**

Context config

Description This command enables the context to configure 802.1ag Connectivity Fault Management (CFM) parameters.

domain

Syntax **domain** *md-index* [**format** {**dns** | **mac** | **none** | **string**}] [**name** *md-name*] **level** *level*
domain *md-index*
no domain *md-index*

Context config>eth-cfm

Description This command configures CFM domain parameters.

The **dns**, **mac**, and **string** keywords apply to dot1ag. The **none** keyword applies to Y.1731. Using the **none** keyword means that the association command must use the **icc-based** format. A MEP associated with domain format **none** and association format **icc-based** is a Y.1731 MEP; otherwise, the MEP is a dot1ag MEP.

The **no** form of the command removes the MD index parameters from the configuration.

Parameters *md-index* — specifies the Maintenance Domain (MD) index value

Values 1 to 4294967295

format {**dns | mac | none | string**} **—** specifies a value that represents the type (format) of the *md-name*

Values **dns**: specifies the DNS name format

mac: X:X:X:X:X:X-u

X: [0 to FF] hex

u: [0 to 65535] decimal

none: no name specified (the domain represents a Y.1731 MEG, not a dot1ag domain)

string: specifies an ASCII string

Default string

md-name — specifies a generic Maintenance Domain (MD) name

Values 1 to 43 characters

level — specifies the integer identifying the maintenance domain level (MD level). Higher numbers correspond to higher-level maintenance domains (those with the greatest physical reach) with the highest values for customers' CFM packets. Lower numbers correspond to lower-level maintenance domains (those with more limited physical reach) with the lowest values for single bridges or physical links.

Values 0 to 7

association

Syntax **association** *ma-index* [**format** {**icc-based | integer | string | vid | vpn-id**}] **name** *ma-name*
association *ma-index*
no association *ma-index*

Context config>eth-cfm>domain

Description This command configures the Maintenance Association (MA) for the domain.

The **integer**, **string**, **vid**, and **vpn-id** keywords apply to dot1ag MAs. The **icc-based** keyword applies to Y.1731 MEGs, and is only available when the domain format is **none**. A MEP associated with domain format **none** and association format **icc-based** is a Y.1731 MEP; otherwise the MEP is a dot1ag MEP.

Parameters *ma-index* — specifies the MA index value

Values 1 to 4294967295

format {icc-based | integer | string | vid | vpn-id} — specifies a value that represents the type (format) of the *ma-name*

Values **icc-based**: raw ASCII, exactly 13 characters (the association is a Y.1731 MEG, not a dot1ag MA)

integer: 0 to 65535 (integer value 0 means the MA is not attached to a VID)

string: raw ASCII

vid: 0 to 4094

vpn-id: RFC 2685, Virtual Private Networks Identifier

XXX:XXXX where X is a value between 00 and FF
(for example 00164D:AABBCCDD)

Default integer

ma-name — specifies the part of the maintenance association identifier that is unique within the maintenance domain name

Values 1 to 45 characters

bridge-identifier

Syntax **[no] bridge-identifier** *bridge-id*

Context config>eth-cfm>domain>association

Description This command configures the service ID for the domain association. The *bridge-id* should be configured to match the *service-id* of the service where MEPs for this association will be created. For example, for Epipe service-id 2, set the bridge-id to 2. There is no verification that the service with a matching *service-id* exists.

Parameters *bridge-id* — specifies the bridge ID for the domain association

Values 1 to 2147483647

vlan

Syntax **vlan** *vlan-id*
no vlan

Context config>eth-cfm>domain>association>bridge-identifier

Description This command configures the bridge-identifier primary VLAN ID. Note that it is informational only, and no verification is done to ensure that MEPs on this association are on the configured VLAN.

Parameters *vlan-id* — specifies a VLAN ID monitored by MA

Values 0 to 4094

ccm-interval

Syntax **ccm-interval** {**10ms** | **100ms** | **1** | **10** | **60** | **600**}
no ccm-interval

Context config>eth-cfm>domain>association

Description This command configures the CCM transmission interval for all MEPs in the association, in milliseconds and seconds.

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

Default **10 s**

remote-mepid

Syntax [**no**] **remote-mepid** *mep-id*

Context config>eth-cfm>domain>association

Description This command configures the remote maintenance association endpoint MEP identifier.

Parameters *mep-id* — maintenance association endpoint identifier of a remote MEP whose information from the MEP database is to be returned

Values 1 to 8191

cfm-loopback

Syntax **cfm-loopback priority** {**low** | **high**}
no cfm-loopback

Context config>port>ethernet

Description This command enables the port to respond to LBM messages and sets the queuing and scheduling conditions for handling CFM LBM frames. The user selects the desired QoS treatment by enabling the CFM loopback and including high or low priority with the **high** or **low** keyword. The queue parameters and scheduler mappings associated with the **high** and **low** keywords are preconfigured and cannot be altered by the user. These parameters and mappings have the following settings:

- for network egress, where profiled scheduling is the choice of scheduling:
  - → **high-priority**: either cir = port_speed, which applies to all frames that are scheduled via an in-profile scheduler, or round-robin (RR) for all other (network egress queue) frames that are in-profile
  - → **low-priority**: either cir = 0, pir = port_speed, which applies to all frames that are scheduled as out-of-profile, or RR for all other frames that are out-of-profile

- for network egress or access egress, where 4-priority scheduling is enabled:
  - → **high-priority**: either cir = port_speed, which applies to all frames that are scheduled via an expedited in-profile scheduler, or RR for all other (network egress queue) frames that reside in expedited queues and are in an in-profile state
  - → **low-priority**: either cir = 0, pir = port_speed, which applies to all frames that are scheduled via a best effort out-of-profile scheduler, or RR for all other frames that reside in best-effort queues and are in an out-of-profile state

The **no** form of the command disables the handling of CFM loopback frames.

Default **no cfm-loopback**

Parameters **low —** sets the queue parameters and scheduler mappings, as described above

high — sets the queue parameters and scheduler mappings, as described above

eth-cfm

Syntax **eth-cfm**

Context config>service>epipe>sap
config>service>epipe>spoke-sdp

Description This command enables the context to configure ETH-CFM parameters.

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up** | **down**}]
no mep *mep-id* **domain** *md-index* **association** *ma-index*

Context config>service>epipe>sap>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm

Description This command provisions an 802.1ag or a Y.1731 maintenance association endpoint (MEP).

The 7705 SAR supports Up and Down MEPs on Ethernet SAPs (802.1ag and Y.1731), and Down MEPs on Ethernet spoke SDPs (802.1ag only).

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

Parameters *mep-id* — specifies the maintenance association endpoint identifier

Values 1 to 81921

md-index — specifies the maintenance domain (MD) index value

Values 1 to 4294967295

ma-index — specifies the MA index value

Values 1 to 4294967295

up | down — specifies the direction in which the maintenance association (MEP) faces on the bridge port (**up** sends Continuity Check messages (CCMs) towards the fabric, **down** sends CCMs towards the egress port or line)

ais-enable

Syntax [**no**] **ais-enable**

Context config>service>epipe>sap>eth-cfm>mep

Description This command enables the generation and the reception of AIS messages and applies to Y.1731 SAP MEPs only.

Default **disabled**

client-meg-level

Syntax **client-meg-level** [*level* [*level* **...**]]
no client-meg-level

Context config>service>epipe>sap>eth-cfm>mep>ais-enable

Description This command configures the client Maintenance Entity Group (MEG) level(s) to use for AIS message generation. Up to seven levels can be provisioned, with the restriction that the client (remote) MEG level must be higher than the local MEG level.

Parameters *level* — specifies the client MEG level

Values 1 to 7

Default 1

interval

Syntax **interval** {**1** | **60**}
no interval

Context config>service>epipe>sap>eth-cfm>mep>ais-enable

Description This command specifies the transmission interval of AIS messages in seconds.

Parameters **1 | 60 —** the transmission interval of AIS messages in seconds

Default 1

priority

Syntax **priority** *priority-value*
no priority

Context config>service>epipe>sap>eth-cfm>mep>ais-enable

Description This command specifies the priority of AIS messages originated by the MEP, which is used for priority-mapping OAM frames.

Parameters *priority-value* — specifies the priority value of the AIS messages originated by the node

Values 0 to 7

Default 7

ccm-enable

Syntax [**no**] **ccm-enable**

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax **ccm-ltm-priority** *priority*
no ccm-ltm-priority

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command specifies the priority value for Continuity Check messages (CCMs) and linktrace messages (LTMs) transmitted by the MEP.

The default priority is 7, which means that CCM frames map to the NC forwarding class by default.

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

Default **7**

Parameters *priority* — specifies the priority of CCM and LTM messages

Values 0 to 7

dual-ended-loss-test-enable

Syntax [**no**] **dual-ended-loss-test-enable**

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command enables dual-ended loss measurement testing on a MEP. When enabled, the test runs in the background.

CCM must be enabled before the dual-ended loss measurement test can be enabled.

The dual-ended and single-ended loss measurement tests are mutually exclusive tests. When the dual-ended loss measurement test is enabled, the single-ended test is not available.

The **no** form of the command disables the dual-ended loss measurement test.

This command applies only to Y.1731 MEPs.

Default **enabled**

alarm-threshold

Syntax **alarm-threshold** *percentage*
no alarm-threshold

Context config>service>epipe>sap>eth-cfm>mep>dual-ended-loss-test-enable
config>service>epipe>spoke-sdp>eth-cfm>mep>dual-ended-loss-test-enable

Description This command specifies the alarm threshold ratio for frame loss measurement, where *percentage* is defined as (the total number of Tx frames) divided by (the total number of frames dropped) expressed as a percentage. When the alarm threshold is reached, an alarm is raised.

The **no** form of the command removes the priority value from the configuration. Setting the percentage to 0.00 is equivalent to using the **no** form of the command.

Parameters *percentage* — 0.00 to 100.00, adjustable in 0.01% increments

Default 0.25

alarm-clear-threshold

Syntax **alarm-clear-threshold** *percentage*
[**no**] **alarm-clear-threshold**

Context config>service>epipe>sap>eth-cfm>mep>dual-ended-loss-test-enable

Description This command configures a clearing alarm threshold for frame loss measurement, where *percentage* is defined as (the total number of Tx frames) divided by (the total number of frames dropped) expressed as a percentage.

If a dual-ended-loss alarm is outstanding and the alarm-clear-threshold is configured to a non-zero value, the dual-ended-loss clear alarm will not be raised until the dual-ended-loss ratio drops below the alarm-clear-threshold. If the alarm-clear-threshold is configured to 0, the dual-ended-loss clear alarm is raised immediately when the dual-ended-loss ratio drops below the alarm threshold.

This functionality prevents too many alarms from being generated if the loss ratio is toggling above and below the alarm threshold.

The alarm-clear-threshold cannot be greater than the alarm-threshold.

Setting the percentage to 0 means that no alarm-clear-threshold is configured; clear alarm traps will continue to be sent when the loss ratio is no longer above the alarm threshold. This is equivalent to using the **no** form of the command.

Parameters *percentage* — 0.00 to 100.00, adjustable in 0.01% increments

Default 0.00

eth-test-enable

Syntax [**no**] **eth-test-enable**

Context config>service>epipe>sap>eth-cfm>mep

Description This command enables an Ethernet (signal) test (ETH-Test) on a MEP. When enabled, the test runs in the background. This command applies to Y.1731 SAP MEPs only.

For this test, operators must configure ETH-Test parameters on both sender and receiver nodes. The ETH-Test can then be run using the following OAM command:

```
oam eth-cfm eth-test mac-address mep mep-id domain md-index
association ma-index [priority priority] [data-length data-
length]
```

A check is done on the provisioning and the test commands to ensure that the MEP is a Y.1731 MEP. If the MEP is not a Y.1731 MEP, the operation fails and an error message in the CLI and SNMP will indicate the problem. A Y.1731 MEP has domain format **none** and association format **icc-based**.

The **no** form of the command disables the ETH-Test on a MEP.

Default **enabled**

bit-error-threshold

Syntax **bit-error-threshold** *bit-errors*

Context config>service>epipe>sap>eth-cfm>mep>eth-test-enable

Description This command configures a threshold for raising SNMP traps for one-way CFM tests.

For bit-error-threshold tests, test results are available only at the destination node. In order for the network management system to collect the results, SNMP traps need to be raised. This threshold is used to control when to raise a trap. When the number of bit errors reaches the threshold, an SNMP trap is raised.

Configuring a threshold value of 0 will cause the node to raise an SNMP trap for every one-way test it receives.

Parameters *bit-errors* — the bit-error threshold

Values 0 to 11840

Default 1

test-pattern

Syntax [**no**] **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]

Context config>service>epipe>sap>eth-cfm>mep>eth-test-enable
config>service>epipe>spoke-sdp>eth-cfm>mep>eth-test-enable

Description This command configures the test pattern for ETH-Test frames.

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

Parameters **all-zeros | all-ones —** specifies to use all zeros or all ones in the test pattern

Default all-zeros

crc-enable — specifies to generate a CRC checksum

low-priority-defect

Syntax **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}

Context config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep

Description This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default **remErrXcon**

Parameters **allDef** — DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM

macRemErrXcon — DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM

remErrXcon — only DefRemoteCCM, DefErrorCCM, and DefXconCCM

errXcon — only DefErrorCCM and DefXconCCM

xcon — only DefXconCCM

noXcon — no defects DefXcon or lower are to be reported

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>service>epipe>sap>eth-cfm>mep

Description This command configures a threshold for raising SNMP traps for one-way CFM tests.

For one-way-delay-threshold tests, test results are available only at the destination node. In order for the network management system to collect the results, SNMP traps need to be raised. This threshold is used to control when to raise a trap. When the delay time reaches the threshold, an SNMP trap is raised.

Configuring a threshold value of 0 will cause the node to raise an SNMP trap for every one-way test it receives.

Parameters *seconds* — the delay time threshold value

Values 0 to 600

Default 3

Configure SAA (Service Assurance Agent) Commands

saa

Syntax **saa**

Context config

Description This command creates the context to configure the SAA tests.

test

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

Context config>saa

Description This command identifies a test and creates or modifies 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 must be shut down before it can be modified or removed from the configuration.

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

Parameters *test-name* — identifies the SAA test name to be created or edited

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"

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

Default **no description**

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.

jitter-event

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

Context config>saa>test

Description This command 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 generated 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.

Values 0 to 2147483 ms

Default 0

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.

Values 0 to 2147483 ms

Default 0

direction — specifies the direction for OAM ping responses received for an OAM ping test run

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

outbound — monitors the jitter value calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run

roundtrip — monitors the jitter value 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 This command 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.

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.

Values 0 to 2147483647 ms

Default 0

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.

Values 0 to 2147483647 ms

Default 0

direction — specifies the direction for OAM ping responses received for an OAM ping test run

Values **inbound** — monitors the latency value calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run

outbound — monitors the latency value calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run

roundtrip — monitors the latency value 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 This command specifies that at the termination of an SAA test run, 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.

Values 0 to 2147483647 packets

Default 0

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.

Values 0 to 2147483647 packets

Default 0

direction — specifies the direction for OAM ping responses received for an OAM ping test run

Values **inbound** — monitors the loss value calculated for the inbound, one-way, OAM ping responses received for an OAM ping test run

outbound — monitors the loss value calculated for the outbound, one-way, OAM ping requests sent for an OAM ping test run

roundtrip — monitors the loss value calculated for the round-trip, two-way, OAM ping requests and replies for an OAM ping test run

Default roundtrip

type

Syntax [**no**] **type**

Context config>saa>test

Description This command creates the context to provide the test type for the named test. Only a single test type can be configured.

A test can only be modified while the test is in shutdown 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* [**source-mac** *ieee-address*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**ttl** *vc-label-ttl*] [**count** *send-count*] [**send-control**] [**return-control**] [**timeout** *timeout*] [**interval** *interval*]

Context oam
config>saa>test>type

Description This ping utility determines the IP connectivity to a CPE within a specified VPLS service.

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

Values 1 to 2147483647

destination *ip-address* **—** specifies the IP address to be used as the destination for performing an OAM ping operation

source *ip-address* **—** specifies an unused IP address in the same network that is associated with the VPLS

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

Default out

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

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

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

Default be

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

Values 1 to 255

Default 255

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

Values 1 to 255

Default 1

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 that the MAC OAM reply to a data plane MAC OAM request is sent using the control plane instead of the data plane

Default MAC OAM reply sent using the data plane

timeout — specifies the maximum amount of time, in seconds, that the router will wait for a message reply after sending the message request. The **timeout** parameter overrides the default timeout value. 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** value.

Values 1 to 10

Default 5

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

Values 1 to 10

Default 1

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**}]] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context config>saa>test>type

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

Parameters *mac-address* — specifies a unicast destination MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the MD index

Values 1 to 4294967295

ma-index — specifies the MA index

Values 1 to 4294967295

ttl-value — specifies the number of hops to use in a linktrace test

Values 0 to 255

fc-name — specifies the forwarding class for CFM test traffic. The *fc-name* is mapped to the dot1p priority that is set in the CFM frame forwarding class. See Table 7 for the Dot1p Priority-to-FC mapping.

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

Default nc

profile {in | out} — specifies the profile state for CFM test traffic; this parameter is not used

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

Values 1 to 10

Default 1

timeout — specifies the maximum amount of time, in seconds, that the router will wait for a message reply after sending the message request. The **timeout** parameter overrides the default timeout value. 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** value.

Values 1 to 10

Default 5

interval — specifies the minimum amount of time, in seconds, that must expire before the next message request is sent. The **interval** parameter is used to override the default request message send interval. If the **interval** is set to 1 s, and the **timeout** value is set to 10 s, then the maximum time between message requests is 10 s and the minimum is 1 s. The **timeout** value must be less than the **interval** value.

Values 1 to 10

Default 5

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**}]] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context config>saa>test>type

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

Parameters *mac-address* — specifies a unicast destination MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the MD index

Values 1 to 4294967295

ma-index — specifies the MA index

Values 1 to 4294967295

data-size — specifies the packet size in bytes, expressed as a decimal integer

Values 0 to 1500

Default 0

fc-name — specifies the forwarding class for CFM test traffic. The *fc-name* is mapped to the dot1p priority that is set in the CFM frame forwarding class. See Table 7 for the Dot1p Priority-to-FC mapping.

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

Default nc

profile {in | out} — specifies the profile state for CFM test traffic; this parameter is not used

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

Values 1 to 100

Default 1

timeout — specifies the maximum amount of time, in seconds, that the router will wait for a message reply after sending the message request. The **timeout** parameter overrides the default timeout value. 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** value.

Values 1 to 10

Default 5

interval — specifies the minimum amount of time, in seconds, that must expire before the next message request is sent. The **interval** parameter is used to override the default request message send interval. If the **interval** is set to 1 s, and the **timeout** value is set to 10 s, then the maximum time between message requests is 10 s and the minimum is 1 s. The **timeout** value must be less than the **interval** value.

Values 1 to 10

Default 5

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**}]] [**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.

Parameters *mac-address* — specifies a unicast MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number

mep-id — specifies the target MEP ID

Values 1 to 8191

md-index — specifies the MD index

Values 1 to 4294967295

ma-index — specifies the MA index

Values 1 to 4294967295

fc-name — specifies the forwarding class for CFM test traffic. The *fc-name* is mapped to the dot1p priority that is set in the CFM frame forwarding class. See Table 7 for the Dot1p Priority-to-FC mapping.

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

Default nc

profile {in | out} — specifies the profile state for CFM test traffic; this parameter is not used

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

Values 1 to 100

Default 1

timeout — specifies the maximum amount of time, in seconds, that the router will wait for a message reply after sending the message request. The **timeout** parameter overrides the default timeout value. 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** value.

Values 1 to 10

Default 5

interval — specifies the minimum amount of time, in seconds, that must expire before the next message request is sent. The **interval** parameter is used to override the default request message send interval. If the interval is set to 1 s, and the timeout value is set to 10 s, then the maximum time between message requests is 10 s and the minimum is 1 s. The **timeout** value must be less than the **interval** value.

Values 1 to 10

Default 5

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*] [**interval** *seconds*] [{**next-hop** *ip-address*} | {**interface** *interface-name*} | **bypass-routing**] [**count** *requests*] [**do-not-fragment**] [**router** *router-instance*] [**timeout** *timeout*] [**fc** *fc-name* [**profile** {**in** | **out**}]]

Context config>saa>test>type

Description This command configures an ICMP ping test.

Parameters *ip-address* — identifies the far-end IP address to which to send the **icmp-ping** request message in dotted-decimal notation

Values ipv4-address: a.b.c.d

dns-name — identifies the DNS name of the far-end device to which to send the **icmp-ping** request message, expressed as a character string

Values 63 characters maximum

rapid — specifies that packets will be generated as fast as possible instead of the default 1 per second

detail — displays detailed information

time-to-live — specifies the TTL value for the MPLS label, expressed as a decimal integer

Values 1 to 128

Default 64

type-of-service — specifies the service type

Values 0 to 255

Default 0

bytes **—** specifies the request packet size in bytes, expressed as a decimal integer

Values 0 to 16384

Default 56

pattern — specifies the pattern that will be used to fill the date portion in a ping packet. If no pattern is specified, position information will be filled instead.

Values 0 to 65535

source *ip-address* **—** specifies the IP address to be used

Values ipv4-address: a.b.c.d

seconds — defines the minimum amount of time, expressed as a decimal integer, that must expire before the next message request is sent

This parameter is used to override the default request message send interval. If the **interval** is set to 1 s, and the **timeout** value is set to 10 s, then the maximum time between message requests is 10 s and the minimum is 1 s. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Values 1 to 10000

Default 1

next-hop *ip-address* **—** displays only the static routes with the specified next-hop IP address

Values ipv4-address: a.b.c.d (host bits must be 0)

interface-name — specifies the name of an IP 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

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

Values 1 to 100000

Default 5

do-not-fragment — sets the DF (Do not fragment) bit in the ICMP ping packet

router-instance — specifies the router name or service ID

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

Default Base

timeout — specifies 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.

This value is used to override the default timeout value.

Values 1 to 10

Default 5

fc-name — indicates 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 control the mapping back to the internal forwarding class used by the far-end 7705 SAR that receives the message request. The egress mappings of the egress network interface on the far-end router control the forwarding class markings on the return reply message. The LSP-EXP mappings on the receive network interface control the mapping of the message reply back at the originating SAR.

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

Default nc

profile {**in** | **out**} **—** specifies the profile state of the MPLS echo request encapsulation

Default in

icmp-trace

Syntax **icmp-trace** [*ip-address* | *dns-name*] [**ttl** *time-to-live*] [**wait** *milli-seconds*] [**source** *ip-address*] [**tos** *type-of-service*] [**router** *router-instance*]

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 **icmp-trace** request message in dotted-decimal notation

Values ipv4-address: a.b.c.d

dns-name — the DNS name of the far-end device to which to send the **icmp-trace** request message, expressed as a character string

Values 63 characters maximum

time-to-live — the TTL value for the MPLS label, expressed as a decimal integer

Values 1 to 255

milli-seconds — the time, in milliseconds, to wait for a response to a probe, expressed as a decimal integer

Values 1 to 60000

Default 5000

source *ip-address* — specifies the IP address to be used

Values ipv4-address: a.b.c.d

type-of-service — specifies the service type

Values 0 to 255

router-instance — specifies the router name or service ID

Values router-name: Base, management

service-id: 1 to 2147483647

Default Base

lsp-ping

Syntax **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**]

Context oam
config>saa>test>type

Description This command performs in-band LSP connectivity tests using the protocol and data structures defined in 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 **detail** parameter is available only from the **oam** context.

Parameters *lsp-name* — specifies a unique LSP name, up to 32 characters in length

path-name — specifies the name for the LSP path, up to 32 characters in length

ip-prefix/mask — specifies the address prefix and subnet mask of the destination node

Values ipv4-address: a.b.c.d
mask: value must be 32

fc-name — indicates 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 control the mapping back to the internal forwarding class used by the far-end 7705 SAR that receives the message request. The egress mappings of the egress network interface on the far-end 7705 SAR control the forwarding class markings on the return reply message.

The LSP-EXP mappings on the receive network interface control the mapping of the message reply back at the originating 7705 SAR.

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

Default be

profile {**in** | **out**} **—** specifies the profile state of the MPLS echo request encapsulation

Default out

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

Values 80, and 85 to 1500 — prefix-specified ping
92, and 97 to 1500 — LSP name-specified ping

Default 80 — prefix-specified ping
92 — LSP name-specified ping
The system sends the minimum packet size, depending on the type of LSP. No padding is added.

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

Values 1 to 255

Default 255

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 time out or receive a reply before the next message request is sent. The message interval value must be expired before the next message request is sent.

Values 1 to 100

Default 1

timeout — specifies 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.

This value is used to override the default timeout value.

Values 1 to 10

Default 5

interval — specifies 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.

This parameter is used to override the default request message send interval.

Values 1 to 10

Default 1

path-destination *ip-address* — specifies the destination IP address

Values ipv4-address: a.b.c.d (host bits must be 0)

if-name — specifies the name of an IP interface. The name must already exist in the **config>router>interface** context.

next-hop *ip-address* — displays only the static routes with the specified next-hop IP address

Values ipv4-address: a.b.c.d (host bits must be 0)

detail — displays detailed information

lsp-trace

Syntax **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**]

Context oam
config>saa>test>type

Description This command displays the hop-by-hop path for an LSP traceroute using the protocol and data structures defined in RFC 4379 *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*.

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 **detail** parameter is available only from the **oam** context.

Parameters *lsp-name* — specifies a unique LSP name, up to 32 characters in length

path-name — specifies the name for the LSP path, up to 32 characters in length

ip-prefix/mask — specifies the address prefix and subnet mask of the destination node

Values ipv4-address: a.b.c.d (host bits must be 0)
mask: 0 to 32

fc-name — indicates 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 control the mapping back to the internal forwarding class used by the far-end 7705 SAR that receives the message request. The egress mappings of the egress network interface on the far-end 7705 SAR control the forwarding class markings on the return reply message.

The LSP-EXP mappings on the receive network interface control the mapping of the message reply back at the originating 7705 SAR.

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

Default be

profile {**in** | **out**} **—** specifies the profile state of the MPLS echo request encapsulation

Values out

no-response-count — specifies the maximum number of consecutive MPLS echo requests, expressed as a decimal integer, that do not receive a reply before the trace operation fails for a given TTL

Values 1 to 255

Default 5

probes-per-hop — specifies the number of OAM requests sent for a particular TTL value, expressed as a decimal integer

Values 1 to 10

Default 1

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

Values 104 to 1500

Default 104 — the system sends the minimum packet size, depending on the type of LSP. No padding is added.

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

Values 1 to 255

Default 1

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

Values 1 to 255

Default 30

timeout — specifies 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.

This value is used to override the default timeout value.

Values 1 to 60

Default 3

interval — specifies 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.

This parameter is used to override the default request message send interval.

Values 1 to 10

Default 1

path-destination *ip-address* — specifies the destination IP address

Values ipv4-address: a.b.c.d (host bits must be 0)

if-name — specifies the name of an IP interface. The name must already exist in the **config>router>interface** context.

next-hop *ip-address* — displays only the static routes with the specified next-hop IP address

Values ipv4-address: a.b.c.d (host bits must be 0)

detail — displays detailed information

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*] [**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 plane.

A MAC ping reply can be sent using the control plane or the data 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 an FDB 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. If the MAC trace originated from a non-zero SHG, the packets will not go out to the same SHG.

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

Values 1 to 2147483647

dst-ieee-address — the destination MAC address for the OAM MAC request

src-ieee-address — the source MAC address from which the OAM MAC request originates. By default, the system MAC address for the chassis is used.

Values Any unicast MAC value

Default The system MAC address

fc-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

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 size packet necessary to send the request is used.

Values 1 to 65535

Default No OAM packet padding

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

Values 1 to 255

Default 255

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

Values 1 to 100

Default 1

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

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

Values 1 to 10

Default 1

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

Values 1 to 10

Default 5

mac-populate

Syntax **mac-populate** *service-id* **mac** *ieee-address* [**flood**] [**age** *seconds*] [**force**] [**target-sap** *sap-id*] [**send-control**]

Context oam

Description This command populates the FDB with an OAM-type MAC entry indicating the node is the egress node for the MAC address, and it optionally floods the OAM MAC association throughout the service. 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 for the MAC address is forwarded to the control plane (CSM). As a result, if the service on the node has neither an FDB nor an egress SAP, then it is not allowed to initiate a **mac-populate** command.

The MAC address that is populated in the FDB in the provider network is given a type OAM, so that it can be treated distinctly from regular dynamically learned or statically configured MACs. 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 the **mac-populate** command forces the MAC in the table to be type OAM in case it already exists as a dynamic, static, or an OAM-induced learned MAC with some other type of 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 FDB 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** comman or with an FDB clear operation.

When a 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-id* — the service ID of the service to diagnose or manage

Values 1 to 2147483647

ieee-address — the MAC address to be populated

flood — sends the OAM MAC populate to all upstream nodes

Default MAC populate only the local FDB

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

Values 1 to 65535

Default the OAM MAC does not age

force — converts the MAC to an OAM MAC even if it currently is another type of MAC

Default do not overwrite type

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 plane; that is, it is associated with the CPU on the router.

When the **target-sap** *sap-id* value is not specified, the MAC is bound to the CSM. 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)

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-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 FDB and optionally floods the OAM MAC removal throughout the service. A **mac-purge** command 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 a packet with the following fields: the Reply Flags is set to 0 (since no reply is expected), and the Reply Mode and Reserved fields are set to 0. The Ethernet header has the source set to the (system) MAC address and 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 that 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 FDB for forwarding, but it is retained in the FDB 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-id* — the service ID of the service to diagnose or manage

Values 1 to 2147483647

ieee-address **—** the MAC address to be purged (all zeros and multicast not allowed)

flood — sends the OAM MAC purge to all upstream nodes

Default MAC purge only the local FDB

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-trace

Syntax **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*]

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 trace operation is modeled after the IP traceroute utility, which uses ICMP echo request and reply packets with increasing TTL values to determine the hop-by-hop route to a destination IP address. The MAC trace command uses Alcatel-Lucent OAM packets with increasing TTL values to determine the hop-by-hop route to a destination MAC.

In a MAC trace, 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 waits for 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. If the MAC ping originated from a non-zero SHG, the packets will not go out to the same SHG.

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

Values 1 to 2147483647

ieee-address — the destination MAC address to be traced (all zeros not allowed)

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

Default be

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 size packet necessary to send the request is used.

Values 1 to 9198

Default no OAM packet padding

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

Values 1 to 255

Default 1

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

Values 1 to 255

Default 4

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

send-count — the number of MAC OAM requests sent for a particular TTL value, expressed as a decimal integer

Values 1 to 100

Default 1

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

Values 1 to 10

Default 1

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

Values 1 to 10

Default 5

sdp-ping

Syntax **sdp-ping** *orig-sdp-id* [**resp-sdp** *resp-sdp-id*] [**fc** *fc-name* [**profile** {**in** | **out**}]] [**size** *octets*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*]

Context oam
config>saa>test>type

Description This command tests SDPs for unidirectional 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 unidirectional 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. Table 8 displays the response messages sorted by precedence.

Table 8: SDP Ping Response Messages

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
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 SDP tunnel encapsulation used to reach the far end — GRE, IP, or MPLS. If *orig-sdp-id* is invalid or administratively down or unavailable for some reason, the SDP echo request message is not sent and an appropriate error message is displayed (once the interval timer expires, **sdp-ping** will attempt to send the next request if required).

Values 1 to 17407

resp-sdp-id — specifies the return SDP-ID to be used by the far-end 7705 SAR for the message reply for round-trip SDP connectivity testing. If *resp-sdp-id* does not exist on the far-end 7705 SAR, terminates on another 7705 SAR different from the originating 7705 SAR, or another issue prevents the far-end 7705 SAR from using *resp-sdp-id*, the SDP echo reply will be sent using generic OAM encapsulaton. The received forwarding class (as mapped on the ingress network interface for the far end) defines the forwarding class encapsulation for the reply message.

This is an optional parameter.

Values 1 to 17407

Default null – use the non-SDP return path for message reply

fc-name — indicates 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 control the mapping back to the internal forwarding class used by the far-end 7705 SAR that receives the message request. The egress mappings of the egress network interface on the far-end 7705 SAR control the forwarding class markings on the return reply message.

The DSCP or LSP-EXP mappings on the receive network interface control the mapping of the message reply back at the originating 7705 SAR. This is displayed in the response message output upon receipt of the message reply.

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

Default be

profile {**in** | **out**} — specifies the profile state of the SDP encapsulation

Default out

octets — the size of the packet 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 SDP, the IP DF (Do not fragment) bit is set. If any segment of the path between the sender and receiver cannot handle the message size, the message is discarded. MPLS LSPs are not expected to fragment the message either, as the message contained in the LSP is not an IP packet.

Values 72 to 1500

Default 40

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

Values 1 to 100

Default 1

timeout — specifies 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.

This value is used to override the default timeout value.

Values 1 to 10

Default 5

interval — specifies 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.

This parameter is used to override the default request message send interval.

Values 1 to 10

Default 1

Special Cases **Single Response Connectivity Tests —** A single response **sdp-ping** test provides detailed test results.

Upon request timeout, message response, request termination, or request error, the local and remote information described in Table 9 will be displayed. Local and remote information is dependent upon SDP-ID existence and reception of reply**.**

Table 9: Single Response Connectivity

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 shut down, 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

Table 9: Single Response Connectivity (Continued)

Field	Description	Values
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 7705 SAR 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 7705 SAR 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 7705 SAR but is not valid for the originating 7705 SAR, Invalid is displayed. When *resp-sdp-id* does not exist on the far-end 7705 SAR, Non-Existent is displayed. When **resp-sdp** is not specified, N/A is displayed.	Admin-Down Admin-Up Invalid Non-Existent N/A
Responding SDP-ID Operational State	The operational state of the far-end SDP-ID associated with the return path for *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

Table 9: Single Response Connectivity (Continued)

Field	Description	Values
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 originator's SDP-ID from the perspective of the remote 7705 SAR 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> oam sdp-ping 10 resp-sdp 22 fc ef
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

==> IP Interface State:     Up
    Actual IP Address:      10.10.10.11     10.10.10.10
    Expected Peer IP:       10.10.10.10     10.10.10.11

    Forwarding Class        ef              ef
    Profile                 Out             Out

Request Result: Sent - Reply Received
RTT: 30ms
```

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 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 toward the average round-trip time.

Multiple Response Round-trip Connectivity Test Sample Output

```
A:router1> oam sdp-ping 6 resp-sdp 101 size 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*] [**count** *send-count*] [**timeout** *timeout*] [**interval** *interval*] [**ttl** *vc-label-ttl*]

Context oam
config>saa>test>type

Description This command configures a virtual circuit connectivity verification (VCCV) ping test. A VCCV ping test checks connectivity of a VLL in-band. 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 data plane and the control plane. The test is in-band, 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. VCCV ping reuses an LSP ping message format and can be used to test a VLL configured over an MPLS, GRE, or IP SDP.

VCCV ping can be initiated on the terminating provider edge (T-PE) router or the switching provider edge (S-PE) router. The 7705 SAR can function as an S-PE or T-PE. If initiated on the S-PE, the **reply-mode** parameter must be used with the **ip-routed** value. The ping from the T-PE can have values or the values can be omitted.

If a VCCV ping is initiated from a T-PE to a neighboring S-PE (one segment only), only the *sdp-id:vc-id* parameter must be used. However, if the ping is across two or more segments, the *sdp-id:vc-id*, **src-ip-address** *ip-addr*, **dst-ip-address** *ip-addr*, **ttl** *vc-label-ttl* and **pw-id** *pw-id* parameters must be used, where:

- the **src-ip-address** is the system IP address of the router preceding the destination router
- the *pw-id* is the VC ID of the last pseudowire segment
- the *vc-label-ttl* must have a value equal to or greater than the number of pseudowire segments

VCCV ping on multi-segment pseudowires require that the control word be enabled in all segments of the VLL. If the control word is not enabled on spoke SDP it will not be signaled peer VCCV CC bits to the far end, consequently VCCV ping cannot be successfully initiated on that specific spoke SDP.

Parameters *sdp-id:vc-id* — identifies the virtual circuit of the pseudowire being tested. The VC ID must exist on the local router and the far-end peer must indicate that it supports VCCV to allow the user to send a **vccv-ping** message.

This is a mandatory parameter.

Values sdp-id: 1 to 17407
vc-id: 1 to 2147483647

src-ip-address *ip-addr* **—** specifies the source IP address

Values ipv4-address: a.b.c.d

dst-ip-address *ip-addr* **—** specifies the destination IP address

Values ipv4-address: a.b.c.d

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*.

Values 0 to 4294967295

reply-mode {**ip-routed** | **control-channel**} **—** specifies the method for sending the reply message to the far-end 7705 SAR

This is a mandatory parameter.

Values **ip-routed** — indicates a reply mode out-of-band using UDP IPv4

control-channel — indicates a reply mode in-band using VCCV control channel

Default control-channel

fc-name **—** indicates 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 control the mapping back to the internal forwarding class used by the far-end 7705 SAR that receives the message request. The egress mappings of the egress network interface on the far-end router control the forwarding class markings on the return reply message. The LSP-EXP mappings on the receive network interface control the mapping of the message reply back at the originating SAR.

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

Default be

profile {**in** | **out**} **—** specifies the profile state of the MPLS echo request encapsulation

Default out

octets — specifies 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.

Values 88 to 9198

Default 88

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

Values 1 to 100

Default 1

timeout — specifies 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.

This value is used to override the default timeout value.

Values 1 to 10

Default 5

interval — specifies 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.

This parameter is used to override the default request message send interval.

Values 1 to 10

Default 1

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.

Values 1 to 255

Sample Output

Ping from T-PE to T-PE:

```
*A:ALU-dut-b_a# oam vccv-ping 1:1 src-ip-address 5.5.5.5 dst-ip-address 3.3.3.3 pw-id
1 ttl 3
VCCV-PING 1:1 88 bytes MPLS payload
Seq=1, reply from 3.3.3.3 via Control Channel
       udp-data-len=32 rtt=10ms rc=3 (EgressRtr)

---- VCCV PING 1:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 10.0ms, avg = 10.0ms, max = 10.0ms, stddev < 10ms
```

Ping from T-PE to S-PE:

```
*A:ALU-dut-b_a# oam vccv-ping 1:1
VCCV-PING 1:1 88 bytes MPLS payload
Seq=1, reply from 4.4.4.4 via Control Channel
       udp-data-len=32 rtt<10ms rc=8 (DSRtrMatchLabel)

---- VCCV PING 1:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min < 10ms, avg < 10ms, max < 10ms, stddev < 10ms


*A:ALU-dut-b_a# oam vccv-ping 1:1 src-ip-address 4.4.4.4 dst-ip-address 5.5.5.5 ttl 2
pw-id 200
VCCV-PING 1:1 88 bytes MPLS payload
Seq=1, reply from 5.5.5.5 via Control Channel
       udp-data-len=32 rtt<10ms rc=8 (DSRtrMatchLabel)

---- VCCV PING 1:1 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min < 10ms, avg < 10ms, max < 10ms, stddev < 10ms
```

Ping from S-PE (on single or multi-segment):

```
*A:ALU-dut-b_a# oam vccv-ping 4:200 reply-mode ip-routed
VCCV-PING 4:200 88 bytes MPLS payload
Seq=1, reply from 5.5.5.5 via IP
       udp-data-len=32 rtt<10ms rc=8 (DSRtrMatchLabel)

---- VCCV PING 4:200 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min < 10ms, avg < 10ms, max < 10ms, stddev < 10ms


*A:ALU-dut-b_a# oam vccv-ping 4:200 reply-mode ip-routed src-ip-address 5.5.5.5 dst-
ip-address 3.3.3.3 ttl 2 pw-id 1
VCCV-PING 4:200 88 bytes MPLS payload
Seq=1, reply from 3.3.3.3 via IP
       udp-data-len=32 rtt<10ms rc=3 (EgressRtr)

---- VCCV PING 4:200 Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min < 10ms, avg < 10ms, max < 10ms, stddev < 10ms
```

vccv-trace

Syntax **vccv-trace** *sdp-id:vc-id* [**size** *octets*] [**min-ttl** *min-vc-label-ttl*] [**max-ttl** *max-vc-label-ttl*] [**max-fail** *no-response-count*] [**probe-count** *probe-count*] [**reply-mode** {**ip-routed** | **control-channel**}] [**timeout** *timeout-value*] [**interval** *interval-value*] [**fc** *fc-name* [**profile** {**in** |**out**}]] [**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 terminating PE (T-PE). VCCV trace 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 TTL values, 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) includes the next-hop S-PE targeted LDP session source address in the Remote PE Address field of the PW FEC TLV. Each S-PE that terminates and processes the message will include the FEC 128 TLV corresponding to the PW segment to its downstream node in the MPLS echo reply message.The source T-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 should be configured accordingly. However, the T-PE or S-PE node will still probe all hops up to **min-ttl** in order to correctly build the FEC of the desired subset of segments.

Parameters *sdp-id:vc-id* — specifies the VC ID of the pseudowire being tested. The VC ID must exist on the local 7705 SAR and the far-end peer must indicate that it supports VCCV to allow the user to send a VCCV ping message.

Values sdp-id : 1 to 17407

vc-id: 1 to 4294967295

octets — specifies 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.

Values 88 to 9198

Default 88

min-vc-label-ttl — specifies 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. The outer label TTL is still set to the default of 255 regardless of the value of the VC label.

Values 1 to 255

Default 1

max-vc-label-tt — specifies 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. The outer label TTL is still set to the default of 255 regardless of the value of the VC label.

Values 1 to 255

Default 8

no-response-count — specifies the maximum number of consecutive VCCV trace echo requests, expressed as a decimal integer, that do not receive a reply before the trace operation fails for a given TTL value.

Values 1 to 255

Default 5

probe-count — specifies the number of VCCV trace echo request messages to send per TTL value

Values 1 to 10

Default 1

reply-mode {**ip-routed** | **control-channel**} — specifies the method for sending the reply message to the far-end 7705 SAR

This is a mandatory parameter.

Values **ip-routed** — indicates a reply mode out-of-band using UDP IPv4

control-channel — indicates a reply mode in-band using the VCCV control channel

Default control-channel

timeout-value — specifies 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 7705 SAR will wait for a message reply after sending the message request. Upon the expiration of message timeout, the requesting 7705 SAR assumes that the message response will not be received. A request timeout message is displayed by the CLI for each message request sent that expires. Any response received after the request times out will be silently discarded.

Values 1 to 60

Default 3

interval-value — specifies 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 s and the timeout value is set to 10 s, then the maximum time between message requests is 10 s and the minimum is 1 s. This depends upon the receipt of a message reply corresponding to the outstanding message request.

Values 1 to 255

Default 1

fc-name — specifies the forwarding class of the VCCV trace echo request encapsulation. 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 control 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 interfac on the far-end router control the forwarding class markings on the return reply message. The LSP-EXP mappings on the receive network interface control the mapping of the message reply back at the originating router.

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

Default be

profile {**in** | **out**} — specifies the profile state of the VCCV trace echo request encapsulation

Default out

detail — displays detailed information

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:ALU2>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)
-----------------------------------------------
*A:ALU2>oam vccv-trace#
```

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*] [**count** *send-count*] [**return-control**] [**timeout** *timeout*] [**interval** *seconds*]

Context config>saa>test>type

Description This command performs a VPRN ping.

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

Values 1 to 2147483647

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

Values ipv4-address: 0.0.0.0 to 255.255.255.255

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

Values 0.0.0.0 to 255.255.255.255

size — the OAM request packet size in octets, expressed as a decimal integer

Values 1 to 9198

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

Values 1 to 255

Default 255

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

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.

Values 1 to 10

Default 1

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

Values 1 to 100

Default 1

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

Values 1 to 100

Default 5

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

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

Default be

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

Default out

Output **Sample Output**

```
A:PE_1# oam vprn-ping 25 source 10.4.128.1 destination 10.16.128.0
Sequence Node-id Reply-Path Size RTT
-------------------------------------------------------------------------------
[Send request Seq. 1.]
1 10.128.0.3:cpm In-Band 100 0ms
-------------------------------------------------------------------------------
...
A:PE_1#
-------------------------------------------------------------------------------
A:PE_1#
```

vprn-trace

Syntax **vprn-trace** *service-id* **source** *ip-address* **destination** *ip-address* [**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** *seconds*]

Context config>saa>test>type

Description This command performs a VPRN trace.

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

Values 1 to 2147483647

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

Values ipv4-address: 0.0.0.0 to 255.255.255.255

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

Values 0.0.0.0 to 255.255.255.255

size — 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

Values 1 to 255

Default 1

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

Values 1 to 255

Default 4

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

Default OAM reply sent using the data plane.

send-count — the number of OAM requests sent for a particular TTL value, expressed as a decimal integer

Values 1 to 10

Default 1

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.

Values 1 to 10

Default 1

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

Values 1 to 10

Default 3

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

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

Default be

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

Default out

Output **Sample Output**

```
A:PE_1# oam vprn-trace 25 source 10.4.128.1 destination 10.16.128.0
TTL Seq Reply Node-id Rcvd-on Reply-Path RTT
-------------------------------------------------------------------------------
[Send request TTL: 1, Seq. 1.]
1 1 1 10.128.0.4 cpm In-Band 0ms
Requestor 10.128.0.1 Route: 0.0.0.0/0
Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
Next Hops: [1] ldp tunnel
Route Targets: [1]: target:65100:1
Responder 10.128.0.4 Route: 10.16.128.0/24
Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
Next Hops: [1] ldp tunnel
Route Targets: [1]: target:65001:100
[Send request TTL: 2, Seq. 1.]
2 1 1 10.128.0.3 cpm In-Band 0ms
Requestor 10.128.0.1 Route: 0.0.0.0/0
Vpn Label: 131071 Metrics 0 Pref 170 Owner bgpVpn
Next Hops: [1] ldp tunnel
Route Targets: [1]: target:65100:1
Responder 10.128.0.3 Route: 10.16.128.0/24
Vpn Label: 0 Metrics 0 Pref 0 Owner local
Next Hops: [1] ifIdx 2 nextHopIp 10.16.128.0
[Send request TTL: 3, Seq. 1.]
[Send request TTL: 4, Seq. 1.]
...
-------------------------------------------------------------------------------
A:PE_1#
```

enable-icmp-vse

Syntax [**no**] **enable-icmp-vse**

Context config>system

Description This command is a global command that enables and disables one-way timestamping of outbound SAA ICMP ping packets. Enabling one-way timestamping on a 7705 SAR node requires **enable-icmp-vse** to be set on both the near-end and far-end nodes. The current status can be seen on the **show>system>information** CLI display.

The **-vse** part of the command means vendor-specific extension.

The **no** form of this command disables one-way timestamping.

Default **no enable-icmp-vse**

OAM SAA Commands

saa

Syntax **saa** *test-name* [**owner** *test-owner*] {**start** | **stop**}

Context oam

Description This command starts or stops an SAA test.

Parameters *test-name* — specifies the name of the SAA test to be run. The test name must already be configured in the **config>saa>test** context.

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 — starts the test. A test cannot be started if the same test is still running or if the test is in a shutdown state. An error message and log event will be generated to indicate a failed attempt to start an SAA test run.

stop — stops a test in progress. A log message will be generated to indicate that an SAA test run has been aborted.

Show Commands

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 dot1ag and Y.1731 association information.

Parameters *ma-index* — specifies the MA index

Values 1 to 4294967295

detail — displays detailed information for the association

Output The following output is an example of eth-cfm association information, and Table 10 describes the fields.

Sample Output

```
*A:ALU-1>show>eth-cfm# association
==============================================================================
Dot1ag CFM Association Table
==============================================================================
Md-index   Ma-index   Name                      CCM-interval Bridge-id
------------------------------------------------------------------------------
1          1          kanata_MA                 10           2
1          2          2                         10           20
==============================================================================
*A:ALU-1>show>eth-cfm#

*A:ALU-1>show>eth-cfm# association detail
-------------------------------------------------------------------------------
Domain 1 Associations:
-------------------------------------------------------------------------------
Md-index          : 1                         Ma-index          : 1
Name Format       : charString                CCM-interval      : 10
Name              : kanata_MA
Bridge-id         : 2                         MHF Creation      : defMHFnone
PrimaryVlan       : 2                         Num Vids          : 0
```

```
-------------------------------------------------------------------------------
Domain 2 Associations:
-------------------------------------------------------------------------------
Md-index          : 2                         Ma-index          : 2
Name Format       : icc-based                 CCM-interval      : 100ms
Name              : 1234567890123
Bridge-id         : 2                         MHF Creation      : defMHFnone
PrimaryVlan       : 2                         Num Vids          : 0
Remote Mep Id     : 2

-------------------------------------------------------------------------------
*A:ALU-1>show>eth-cfm#
```

Table 10: ETH-CFM Association Field Descriptions

Label	Description
Md-index	Displays the MD index
Ma-index	Displays the MA index
Name	Displays the name of the MA
CCM-interval	Displays the CCM interval (in seconds)
Bridge-id	Displays the bridge ID for the MA. The bridge ID is the same value as the service ID of the service to which the MEP belongs.
Name Format	Displays the format for the MA name
MHF Creation	Not applicable
PrimaryVlan	Displays the VLAN ID
Num Vids	Displays the number of VLAN IDs
Remote Mep Id	Displays the MEP identifier for the remote MEP

cfm-stack-table

Syntax **cfm-stack-table**
cfm-stack-table port [*port-id* [**vlan** *vlan-id*]] [**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*]

Context show>eth-cfm

Description This command displays stack-table information.

Parameters *port-id* — displays the bridge port or aggregated port on which MEPs are configured

Values *slot*/*mda*/*port*[.*channel*]

vlan-id — displays the associated VLAN ID

Values 0 to 4094

sdp-id[:*vc-id*] — displays the SDP binding for the bridge

Values *sdp-id* 1 to 17407

vc-id 1 to 4294967295

0...7 — display the MD level of the maintenance point

Values 0 to 7

service-id — displays the CFM stack table information for the specified *service-id*

Values 0 to 2147483647

up | down — displays the direction that the MEP faces on the bridge port

Output The following output is an example of eth-cfm stack table information, and Table 11 describes the fields.

Sample Output

```
*A:ALU-1>show>eth-cfm# cfm-stack-table
===============================================================================
CFM SAP Stack Table
===============================================================================
Sap            Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
1/5/1          5     Down 1          1          1
===============================================================================

===============================================================================
CFM SDP Stack Table
===============================================================================
Sdp            Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
1:11           5     Down 1          1          2      a4:58:ff:00:00:00
===============================================================================

===============================================================================
CFM Virtual Stack Table
===============================================================================
Service        Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALU-1>show>eth-cfm#
```

Table 11: ETH-CFM Stack Table Field Descriptions

Label	Description
Sap	Displays the SAP identifier
Sdp	Displays the spoke SDP identifier
Service	Displays the service identifier
Level	Displays the MD level of the domain
Dir(direction)	Displays the direction of OAMPDU transmission
Md-index	Displays the MD index of the domain
Mep-id	Displays the MEP identifier
Mac-address	Displays the MAC address of the MEP

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 MEP is associated, or 0, if none

Values 1 to 4294967295

ma-index — displays the index to which the MA is associated, or 0, if none

Values 1 to 4294967295

all-associations — displays all associations to the MD

detail — displays detailed domain information

Output The following output is an example of eth-cfm domain information, and Table 12 describes the fields.

Sample Output

```
*A:ALU-1>show>eth-cfm# domain
===============================================================================
CFM Domain Table
===============================================================================
Md-index   Level Name                                          Format
-------------------------------------------------------------------------------
1          5     kanata_MD                                    charString
2          1                                                  none
===============================================================================
```

```
*A:ALU-1>show>eth-cfm# domain detail
===============================================================================
Domain 1
Md-index          : 1                         Level            : 5
Permission        : sendIdNone                MHF Creation     : defMHFnone
Name Format       : charString                Next Ma Index    : 2
Name              : kanata_MD
===============================================================================
Domain 2
Md-index          : 2                         Level            : 1
Permission        : sendIdNone                MHF Creation     : defMHFnone
Name Format       : none                      Next Ma Index    : 1
===============================================================================

*A:ALU-1>show>eth-cfm# domain all-associations
===========================================================================
CFM Association Table
===========================================================================
Md-index   Ma-index   Name                      CCM-interval Bridge-id
---------------------------------------------------------------------------
1          1          kanata_MA                 10           2
2          2          1234567890123             100ms        2
===========================================================================

*A:ALU-1>show>eth-cfm# domain all-associations detail
===============================================================================
Domain 1
Md-index          : 1                         Level            : 5
Permission        : sendIdNone                MHF Creation     : defMHFnone
Name Format       : charString                Next Ma Index    : 2
Name              : kanata_MD
-------------------------------------------------------------------------------
Domain 1 Associations:

Md-index          : 1                         Ma-index         : 1
Name Format       : string                    CCM-interval     : 10
Name              : kanata_MA
Bridge-id         : 2                         MHF Creation     : defMHFnone
PrimaryVlan       : 2                         Num Vids         : 0
Remote Mep Id     : 1

===============================================================================
Domain 2
Md-index          : 2                         Level            : 1
Permission        : sendIdNone                MHF Creation     : defMHFnone
Name Format       : none                      Next Ma Index    : 1
-------------------------------------------------------------------------------
Domain 2 Associations:

Md-index          : 2                         Ma-index         : 2
Name Format       : icc-based                 CCM-interval     : 100ms
Name              : 1234567890123
Bridge-id         : 2                         MHF Creation     : defMHFnone
PrimaryVlan       : 2                         Num Vids         : 0
Remote Mep Id     : 2

===============================================================================
*A:ALU-1>show>eth-cfm#
```

Table 12: ETH-CFM Domain Field Descriptions

Label	Description
Domain	
Md-index	Displays the MD index of the domain
Level	Displays the MD level of the domain
Permission	Not applicable
MHF Creation	Not applicable
Name Format	Displays the format for the MD name
Next Ma Index	Displays the value of the next MA index
Name	Displays the name of the MD
Domain Associations	
Md-index	Displays the MD index of the domain
Ma-index	Displays the MA index of the association
Name Format	Displays the format for the MA name
CCM-interval	Displays the CCM interval (in seconds)
Name	Displays the name of the MA
Bridge-id	Displays the bridge ID for the MA. The bridge ID is the same value as the service ID of the service to which the MEP belongs.
MHF Creation	Not applicable
PrimaryVlan	Displays the VLAN ID configured under the **config>eth-cfm>domain>association>bridge-identifier>vlan** command
Num Vids	Displays the number of VLAN IDs and is always 0
Remote Mep Id	Displays the MEP identifier for the remote MEP

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* **single-ended-loss-test** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **dual-ended-loss-test** [**remote-peer** *mac-address*]

Context show>eth-cfm

Description This command displays information for various Ethernet OAM tests and entities related to MEPs, including:

- MEPs
- loopback
- linktrace
- remote MEPs
- Ethernet signal test
- delay and delay variation measurements (one-way and two-way)
- loss measurements (single-ended and dual-ended)

Parameters *mep-id* — specifies the target MEP ID

Values 1 to 8191

md-index — displays the index of the MD to which the MEP is associated, or 0, if none

Values 1 to 4294967295

ma-index — displays the index of the MA to which the MEP is associated, or 0, if none

Values 1 to 4294967295

mac-address — displays the MAC address of the remote peer MEP

Values xx:xx:xx:xx:xx:xx or
xx-xx-xx-xx-xx-xx,
where xx is a hexadecimal number

loopback — displays loopback information for the specified MEP

linktrace — displays linktrace information for the specified MEP

remote-mepid — displays specified remote *mep-id* information for the specified MEP

all-remote-mepids — displays all remote *mep-id* information for the specified MEP

remote-peer — displays specified remote *mep-id* information for the specified MEP

eth-test-results — displays ETH-Test result information for the specified MEP and remote peer

one-way-delay-test — displays one-way test information for the specified MEP and remote peer

two-way-delay-test — displays two-way test information for the specified MEP and remote peer

single-ended-loss-test — displays single-ended-loss test information for the specified MEP and remote peer

dual-ended-loss-test — displays dual-ended-loss test information for the specified MEP and remote peer

Output The following outputs are examples of Ethernet OAM tests for MEPs:

- MEPs, Loopback, and Linktrace (Sample Output, Table 13)
- Remote MEPs (Sample Output, Table 14)
- ETH-Test results (Sample Output, Table 15)
- Delay measurements (one-way and two-way) (Sample Output (one-way) and Sample Output (two-way), Table 16)
- Loss test (single-ended and dual-ended) (Sample Output (single-ended) and Sample Output (two-way), Table 17)

Sample Output

```
*A:ALU-1>show>eth-cfm# mep 2 domain 1 association 1 loopback linktrace
-------------------------------------------------------------------------------
Mep Information
-------------------------------------------------------------------------------
Md-index           : 2                          Direction         : Down
Ma-index           : 20                         Admin             : Enabled
MepId              : 200                        CCM-Enable        : Enabled
IfIndex            : 46333952                   PrimaryVid        : 200
FngState           : fngReset
LowestDefectPri    : macRemErrXcon              HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:25:ba:30:2e:1f          CcmLtmPriority    : 7
CcmTx                 : 188                     CcmSequenceErr    : 0
DmrRepliesTx       : 0
LmrRepliesTx       : 0                          Dual-Loss Thresh  : 1.20%
Dual-Loss Test     : Enabled                    Dual-Loss AlarmClr: 0.80%
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None

-------------------------------------------------------------------------------
Mep Loopback Information
-------------------------------------------------------------------------------
LbRxReply          : 0                        LbRxBadOrder      : 0
LbRxBadMsdu        : 0                        LbTxReply         : 0
LbSequence         : 1                        LbNextSequence    : 1
LbStatus           : False                    LbResultOk        : False
```

```
DestIsMepId        : False                    DestMepId         : 0
DestMac            : 00:00:00:00:00:00        SendCount         : 0
VlanDropEnable     : True                     VlanPriority      : 7
Data TLV:
    None
-------------------------------------------------------------------------------
Mep Linktrace Message Information
-------------------------------------------------------------------------------
LtRxUnexplained    : 0                        LtNextSequence    : 1
LtStatus           : False                    LtResult          : False
TargIsMepId        : False                    TargMepId         : 0
TargMac            : 00:00:00:00:00:00        TTL               : 64
EgressId           : 00:00:a4:58:ff:00:00:00  SequenceNum       : 1
LtFlags            : useFDBonly
-------------------------------------------------------------------------------
Mep Linktrace Replies
-------------------------------------------------------------------------------
SequenceNum        : 1                        ReceiveOrder      : 1
Ttl                : 63                       Forwarded         : False
LastEgressId       : 00:00:00:21:05:6e:5a:f1  TerminalMep       : True
NextEgressId       : 00:00:00:21:05:4d:a8:b2  Relay             : rlyHit
ChassisIdSubType   : unknown value (0)
ChassisId:
    None
ManAddressDomain:
    None
ManAddress:
    None
IngressMac         : 00:21:05:4d:a8:b2        Ingress Action    : ingOk
IngrPortIdSubType  : unknown value (0)
IngressPortId:
    None
EgressMac          : 00:00:00:00:00:00        Egress Action     : egrNoTlv
EgrPortIdSubType   : unknown value (0)
EgressPortId:
    None
Org Specific TLV:
    None
-------------------------------------------------------------------------------
*A:ALU-1>show>eth-cfm#
```

Table 13: ETH-CFM MEP, Loopback, and Linktrace Field Descriptions

Label	Description
Mep Information	
Md-index	Displays the MD index of the domain
Direction	Displays the direction of OAMPDU transmission
Ma-index	Displays the MA index of the association
Admin	Displays the administrative status of the MEP
MepId	Displays the MEP identifier

Table 13: ETH-CFM MEP, Loopback, and Linktrace Field Descriptions (Continued)

Label	Description
CCM-Enable	Displays the status of the CCM (enabled or disabled)
IfIndex	Displays the index of the interface
PrimaryVid	Displays the identifier of the primary VLAN
FngState	Indicates the different states of the Fault Notification Generator
LowestDefectPri	Displays the lowest priority defect (a configured value) that is allowed to generate a fault alarm
HighestDefect	Identifies the highest defect that is present (for example, if defRDICCM and defXconCCM are present, the highest defect is defXconCCM)
Defect Flags	Displays the number of defect flags
Mac Address	Displays the MAC address of the MEP
CcmLtmPriority	Displays the priority value transmitted in the linktrace messages (LTM)s and CCMs for this MEP. The MEP must be configured on a VLAN.
CcmTx	Displays the number of Continuity Check Messages (CCM) sent The count is taken from the last polling interval (every 10 s)
CcmSequenceErr	Displays the number of CCM errors
Eth-1DM Threshold	Displays the one-way-delay threshold value
DmrRepliesTx	Displays the number of delay measurement replies transmitted
LmrRepliesTx	Displays the number of loss measurement replies transmitted
Dual-Loss-Test	Displays the state of the dual-ended loss test (enabled or disabled)
Dual-Loss Threshold	Displays the alarm threshold for frame loss measurement
Dual-Loss AlarmClr	Displays the clearing alarm threshold for frame loss measurement
Eth-Ais	Displays the state of the ETH-AIS test (enabled or disabled)
Eth-Test	Displays the state of the ETH-Test (enabled or disabled)
Eth-Test dataLength	Displays the data length of the MEP
Eth-Test Threshold	Displays the bit-error threshold setting
Eth-Test Pattern	Displays the test pattern configured for the MEP

Table 13: ETH-CFM MEP, Loopback, and Linktrace Field Descriptions (Continued)

Label	Description
Eth-Test Priority	Displays the priority of frames with ETH-Test information
CcmLastFailure Frame	Displays the frame that caused the last CCM failure
XconCcmFailure Frame	Displays the frame that caused the XconCCMFailure
Mep Loopback Information	
LbRxReply	Displays the number of received loopback (LB) replies
LbRxBadOrder	Displays the number of received loopback messages that are in a bad order
LbRxBadMsdu	Displays the number of loopback replies that have been received with the wrong destination MAC address (MSDU = MAC Service Data Unit)
LbTxReply	Displays the number of loopback replies transmitted out this MEP
LbTxReply (Total)	Displays the total number of LBRs (loopback replies) transmitted from this MEP
LbTxReplyNoTLV	Displays the number of LBRs (loopback replies) transmitted from this MEP with no TLV Because only LBMs with no TLVs are used for throughput testing, the LbTxReply (Total), LbTxReplyNoTLV, and LbTxReplyWithTLV counters can help debug problems if throughput testing is not working
LbTxReplyWithTLV	Displays the number of LBRs (loopback replies) transmitted from this MEP with TLV
LbSequence	Displays the sequence number in the loopback message
LbNextSequence	Displays the next loopback sequence
LbStatus	Displays the loopback status as True or False: True – loopback is in progress False – no loopback is in progress
LbResultOk	Displays the result of the loopback test
DestIsMepId	Identifies whether the destination interface has a MEP-ID (true or false)
DestMepId	Displays the MEP-ID of the destination interface
DestMac	Displays the MAC address of the destination interface

Table 13: ETH-CFM MEP, Loopback, and Linktrace Field Descriptions (Continued)

Label	Description
SendCount	Indicates the number of loopback messages sent
VlanDropEnable	Identifies whether the VLAN drop is enabled (true or false)
VlanPriority	Displays the VLAN priority
Data TLV	Displays the data TLV information
Mep Linktrace Message Information	
LtRxUnexplained	Displays the number of unexplained linktrace messages (LTM) that have been received
LtNextSequence	Displays the sequence number of the next linktrace message
LtStatus	Displays the status of the linktrace
LtResult	Displays the result of the linktrace
TargIsMepId	Identifies whether the target interface has a MEP-ID (true or false)
TargMepId	Displays the MEP-ID of the target interface
TargMac	Displays the MAC address of the target interface
TTL	Displays the TTL value
EgressId	Displays the egress ID of the linktrace message
SequenceNum	Displays the sequence number of the linktrace message
LtFlags	Displays the linktrace flags
Mep Linktrace Replies	
SequenceNum	Displays the sequence number returned by a previous transmit linktrace message, indicating which linktrace message response will be returned
ReceiveOrder	Displays the order in which the linktrace initiator received the linktrace replies
Ttl	Displays the TTL field value for a returned linktrace reply
Forwarded	Indicates whether the linktrace message was forwarded by the responding MEP

Table 13: ETH-CFM MEP, Loopback, and Linktrace Field Descriptions (Continued)

Label	Description
LastEgressId	Displays the last egress identifier returned in the linktrace reply egress identifier TLV of the linktrace reply The last egress identifier identifies the MEP linktrace initiator that initiated, or the linktrace responder that forwarded, the linktrace message for which this linktrace reply is the response This is the same value as the egress identifier TLV of that linktrace message
TerminalMep	Indicates whether the forwarded linktrace message reached a MEP enclosing its MA
NextEgressId	Displays the next egress identifier returned in the linktrace reply egress identifier TLV of the linktrace reply The next egress identifier identifies the linktrace responder that transmitted this linktrace reply and can forward the linktrace message to the next hop This is the same value as the egress identifier TLV of the forwarded linktrace message, if any
Relay	Displays the value returned in the Relay Action field
ChassisIdSubType	Displays the format of the chassis ID returned in the Sender ID TLV of the linktrace reply, if any This value is meaningless if the chassis ID has a length of 0
ChassisId	Displays the chassis ID returned in the Sender ID TLV of the linktrace reply, if any The format is determined by the value of the ChassisIdSubType
ManAddressDomain	Displays the TDomain that identifies the type and format of the related ManAddress, used to access the SNMP agent of the system transmitting the linktrace reply Received in the linktrace reply Sender ID TLV from that system
ManAddress	Displays the TAddress that can be used to access the SNMP agent of the system transmitting the CCM Received in the CCM Sender ID TLV from that system
IngressMac	Displays the MAC address returned in the ingress MAC address field

Table 13: ETH-CFM MEP, Loopback, and Linktrace Field Descriptions (Continued)

Label	Description
Ingress Action	Displays the value returned in the Ingress Action field of the linktrace message
IngressPortIdSubType	Displays the format of the ingress port ID
IngressPortId	Displays the ingress port ID; the format is determined by the value of the IngressPortIdSubType
EgressMac	Displays the MAC address returned in the egress MAC address field
Egress Action	Displays the value returned in the Egress Action field of the linktrace message
EgressPortIdSubType	Displays the format of the egress port ID
EgressPortId	Displays the egress port ID; the format is determined by the value of the EgressPortIDSubType
Org Specific TLV	Displays all organization-specific TLVs returned in the linktrace reply, if any Includes all octets including and following the TLV length field of each TLV, concatenated

Sample Output

```
*A:ALU-1>show eth-cfm mep 1 domain 103 association 99 all-remote-mepids
===============================================================================
Eth-CFM Remote-Mep Table
===============================================================================
R-mepId Rx CC  Rx Rdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
-------------------------------------------------------------------------------
2       True   False  Up       Up     8a:d9:ff:00:00:00 02/17/2009 16:27:48
3       True   False  Up       Up     8a:da:01:01:00:02 02/17/2009 16:27:48
===============================================================================
===============================================================================
*A:ALU-1>
*A:ALU-1>show eth-cfm mep 1 domain 103 association 99 remote-mepid 3
===============================================================================
Eth-CFM Remote-Mep Table
===============================================================================
R-mepId Rx CC  Rx Rdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
-------------------------------------------------------------------------------
3       True   False  Up       Up     8a:da:01:01:00:02 02/17/2009 16:27:48
===============================================================================
*A:ALU-1>
```

Table 14: ETH-CFM MEP Remote MEP Field Descriptions

Label	Description
R-mepId	Displays the remote MEP identifier
Rx CC	Displays the state of received CCMs (True or False): True — CCMs are received False — CCMs are not received
Rx Rdi	Displays the state of received RDIs (True or False): True — RDIs are received False — RDIs are not received
Port-Tlv	Displays the contents of the port status TLV in the CCM (Up, Blocked, or Absent), as defined in the 802.1ag specification
If-Tlv	Displays the contents of the interface status TLV in the CCM (Up, Blocked, or Absent), as defined in the 802.1ag specification
Peer Mac Addr	Displays the MAC address of the peer (remote) entity
CCM status since	Displays the date and time when continuity check messages began to be sent

Sample Output

```
*A:ALU-1>show eth-cfm mep 1 domain 103 association 99 eth-test-results
===============================================================
Eth CFM ETH-Test Result Table
===============================================================
                                Current        Accumulate
                  FrameCount    ErrBits        ErrBits
Peer Mac Addr     ByteCount     CrcErrs        CrcErrs
---------------------------------------------------------------
22:34:56:78:9a:bc 1             0              0
                  100           0              0
32:34:56:78:9a:bc 1             0              0
                  100           0              0
42:34:56:78:9a:bc 1             0              0
                  100           0              0
===============================================================
*A:ALU-1>#
*A:ALU-1>show eth-cfm mep 1 domain 103 association 99 eth-test-results remote-peer
22:34:56:78:9a:bc
===============================================================
Eth CFM ETH-Test Result Table
===============================================================
===============================================================
                                Current        Accumulate
                  FrameCount    ErrBits        ErrBits
Peer Mac Addr     ByteCount     CrcErrs        CrcErrs
---------------------------------------------------------------
22:34:56:78:9a:bc 1             0              0
                  100           0              0
===============================================================
*A:ALU-1>
```

Table 15: ETH-CFM MEP ETH-Test Field Descriptions

Label	Description
Peer Mac Addr	Displays the MAC address of the peer (remote) entity
FrameCount	Displays the number of test frames sent between the MEP and the peer entity
ByteCount	Displays the number of bytes sent between the MEP and the peer entity
Current ErrBits	Displays the number of bit errors in the current test
Current CrcErrs	Displays the number of CRC errors in the current test
Accumulate ErrBits	Displays the accumulated number of bit errors in the current test
Accumulate CrcErrs	Displays the accumulated number of CRC errors in the current test

Sample Output (one-way)

```
*A:ALU-1>show eth-cfm mep 1 domain 103 association 99 one-way-delay-test
===============================================================================
Eth CFM One-way Delay Test Result Table
===============================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-------------------------------------------------------------------------------
8a:d8:01:01:00:01     759606              2840
aa:bb:cc:dd:ee:ff     760256              760256
===============================================================================
*A:ALU-1>
*A:ALU-1>show eth-cfm mep 1 domain 103 association 99 one-way-delay-test remote-peer
8a:d8:01:01:00:01
===============================================================================
Eth CFM One-way Delay Test Result Table
===============================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-------------------------------------------------------------------------------
8a:d8:01:01:00:01     759606              2840
===============================================================================
*A:ALU-1>
```

Sample Output (two-way)

```
*A:ALU-1>show eth-cfm mep 2 domain 103 association 99 two-way-delay-test
===============================================================================
Eth CFM Two-way Delay Test Result Table
===============================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-------------------------------------------------------------------------------
00:16:4d:54:49:db     10190               13710
===============================================================================
*A:ALU-1>
*A:ALU-1>show eth-cfm mep 2 domain 103 association 99 two-way-delay-test remote-peer
00:16:4D:54:49:DB
===============================================================================
Eth CFM Two-way Delay Test Result Table
===============================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-------------------------------------------------------------------------------
00:16:4d:54:49:db     10190               13710
===============================================================================
*A:ALU-1>
```

Table 16: ETH-CFM MEP Delay Measurement Test Field Descriptions

Label	Description
Peer Mac Addr	Displays the MAC address of the peer (remote) entity
Delay (us)	Displays the measured delay (in microseconds) for the DM test
Delay Variation (us)	Displays the measured delay variation (in microseconds) for the DV test

Sample Output (single-ended)

```
*A:ALU-1>show eth-cfm mep 1 domain  1 association 1  single-ended-loss-test remote-
peer 00:1a:f0:00:00:01
===============================================================================
Eth CFM Single-Ended Test Result Table
===============================================================================
Far-End Mac Addr:        00:1a:f0:00:00:00     Duration (sec): 5

Latest Frame Counters    In Previous LMR       In Current LMR      Delta
  TxLocal     :          123456                123466               10
  RxFarEnd    :          123450                123460               10
  TxFarEnd    :          123450                123460               10
  RxLocal     :          123456                123465               9


Accumulated Frames                  Near-End             Far-End
   Total Tx       :                 30                   36
   Total Rx       :                 35                   30
   Total Loss     :                 1                    0
   Loss Ratio(%)  :                 2.78                 0.00
===============================================================================
*A:ALU-1>
```

Sample Output (dual-ended)

```
*A:ALU-1>show eth-cfm mep 1 domain 1 association 1 dual-ended-loss-test remote-peer
00:1a:f0:00:00:01
===============================================================================
Eth CFM Dual-Ended Test Result
===============================================================================
Far-End Mac Addr:        00:1a:f0:00:00:01     Duration (sec):    21347
                                               CcmRxCount    :    60632

Latest Frame Counters    In Previous CCM       In Current CCM     Delta
   TxLocal        :      3999                  4000               1
   RxFarEnd       :      3999                  4000               1
   TxFarEnd       :      0                     0                  0
   RxLocal        :      0                     0                  0

Accumulated Frames       Near-End             Far-End
   Total Tx       :      5066117155            741
   Total Rx       :      0                     6720979
   Total Loss     :      741                   5059396176
   Loss Ratio(%)  :      100.00                99.86
===============================================================================
*A:ALU-1>
```

Table 17: ETH-CFM MEP Loss Measurement Test Field Descriptions

Label	Description
Far-End Mac Addr	Displays the MAC address of the far-end (remote) router
Duration(sec)	Displays the duration that the current test has been running Reset via the **clear>eth-cfm>dual-ended-loss-test** command
CCMRxCount	Displays the total number of received CCMs
Latest Frame Counters	Indicates that the number of frames counted are the latest values: • For single-ended tests — the values are for the previous LMR, the current LMR, and the difference between them • For dual-ended tests — the values are the previous CCM, the current CCM, and the difference between them
TxLocal	Displays the latest number of frames transmitted from the local router
RxFarEnd	Displays the latest number of frames received at the remote router
TxFarEnd	Displays the latest number of frames transmitted from the remote router
RxLocal	Displays the latest number of frames received by the local router
Accumulated Frames	Indicates that the frame counter values under this heading are the accumulated values for the near-end (local) and far-end (remote) routers
Total Tx	Displays the total number of frames transmitted during the test
Total Rx	Displays the total number of frames received during the test
Total Loss	Displays the total number of frames lost during the test
Loss Ratio (%)	Displays the loss ratio, defined as follows: • Loss Ratio (NE) = Total Loss (NE) ÷ Total Tx (FE) x 100% Example (single-ended): • NE loss ratio = (1 ÷ 36) x 100% = 2.78% • FE loss ratio = (0 ÷ 30) x 100% = 0.00% Example (dual-ended): • NE loss ratio = (741 ÷ 741) x 100% = 100% • FE loss ratio = (5059396176 ÷ 5066117155) x 100% = 99.86%

saa

Syntax **saa** [*test-name* [**owner** *test-owner*]]

Context show>saa

Description This command displays information about the SAA test.

If no specific test is specified, a summary of all configured tests is displayed.

If a 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* — specifies the SAA test to display. The test name must already be configured in the **config>saa>test** context.

test-owner — specifies the owner of an SAA operation, up to 32 characters in length

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

Output The following output is an example of SAA test result information, and Table 18 describes the fields.

Sample Output

The following displays an SAA test result:

```
*A:ALU-3>config>saa>test$ show saa

===============================================================================
SAA Test Information
===============================================================================
Test name                    : test5
Owner name                   : reuben
Administrative status        : Enabled
Test type                    : sdp-ping 600 resp-sdp 700 fc "nc" count 50
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    50         None       Never               None
            Falling   50         10         04/23/2008 22:29:40 1
```

```
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    8          None       Never                None
            Falling   8          0          04/23/2008 22:30:30 1

===============================================================================
*A:ALU-3>config>saa>test$
```

Table 18: SAA Field Descriptions

Label	Description
Test name	Displays the name of the test
Owner name	Displays the test owner's name
Administrative status	Indicates the administrative state of the test – enabled or disabled
Test type	Identifies the type of test configured
Test runs since last clear	Indicates the total number of tests performed since the last time the tests were cleared
Number of failed tests run	Specifies the total number of tests that failed
Last test result	Indicates the result of the last test 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

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:ALU-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:ALU-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:ALU-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:ALU-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:ALU-48#
```

Clear Commands

saa

Syntax **saa-test** [*test-name* [**owner** *test-owner*]]

Context clear

Description This command clears the SAA results for the specified test and the history for the test. If the test name is omitted, all the results for all tests are cleared.

Parameters *test-name* — specifies the SAA test to clear. The test name must already be configured in the **config>saa>test** context.

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"

dual-ended-loss-test

Syntax **dual-ended-loss-test mep** *mep-id* **domain** *md-index* **association** *ma-index*

Context clear>eth-cfm

Description This command clears the accumulated frame counters during a dual-ended loss measurement (LM) test.

The LM counters are reset when a MEP on the datapath is created or deleted automatically by the OS for network or configuration reasons. Some of the reasons for creating or deleting a MEP are as follows, excluding the general functions of manually creating or deleting a MEP:

- for SAPs
  - → changing the ccm-ltm-priority using the CLI or SNMP
  - → changing the ccm-interval using the CLI or SNMP
  - → changing the SAP egress QoS policy
  - → changes to the SAP state (due to, for example, moving (bouncing) ports, link loss forwarding (LLF), or network changes that require recreation of flows)
- for spoke SDPs
  - → changing the vc type on the spoke SDP
  - → changing the vc vc-tag on the spoke SDP
  - → changing the vc etype on the spoke SDP
  - → change to the spoke SDP state due to network conditions

Note: The **clear>dual-ended loss-test** command only resets the "Accumulated Frames During the Test" results for both the far end and near end. The frame counters for aggregated results are not reset. See Sample Output - before less than two CCMs.

Parameters *mep-id* — specifies the target MEP ID

Values 1 to 8191

md-index — displays the index of the MD to which the MEP is associated, or 0, if none

Values 1 to 4294967295

ma-index — displays the index of the MA to which the MEP is associated, or 0, if none

Values 1 to 4294967295

Output The following outputs show sample displays after issuing a **show>eth-cfm>.....>dual-ended-loss-test** command:

- before receiving two CCMs after issuing the **clear** command
- after receiving two or more CCMs after issuing the **clear** command

Sample Output - before less than two CCMs

```
================================================================================
Eth CFM Dual-Ended Test Result
================================================================================
Far-End Mac Addr  :    00:1a:f0:69:d4:a6       Duration (sec)  : 0

Latest Frame Counters    In Previous CCM       In Current CCM     Delta
   TxLocal        :      0                     0                  0
   RxFarEnd       :      0                     0                  0
   TxFarEnd       :      0                     0                  0
   RxLocal        :      0                     0                  0

Accumulated Frames During Test               Near-End           Far-End
   Total Tx       :                          0                  0
   Total Rx       :                          0                  0
   Total Loss     :                          0                  0
   Loss Ratio(%)  :                          0.00               0.00

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

Sample Output - after two or more CCMs

In the display below, counters that have been cleared and restarted are shown in **bold**.

```
===============================================================================
Eth CFM Dual-Ended Test Result
===============================================================================
Far-End Mac Addr  :    00:1a:f0:69:d4:a6       Duration (sec)  : 2

Latest Frame Counters    In Previous CCM       In Current CCM     Delta
   TxLocal        :      123556                123566             10
   RxFarEnd       :      123550                123560             10
   TxFarEnd       :      123550                123560             10
   RxLocal        :      123556                123566             10

Accumulated Frames During Test               Near-End           Far-End
   Total Tx       :                          10                 10
   Total Rx       :                          10                 10
   Total Loss     :                          0                  0
   Loss Ratio(%)  :                          0.00               0.00

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

Debug Commands

oam

Syntax [**no**] **oam**

Context debug

Description This command enables or disables debugging for OAM.

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 the direction for the LSP ping debugging: transmit, receive, or both transmit and receive

raw | detail — displays output for the debug mode

Tools

Tools Command Reference

Command Hierarchies

- Tools Dump Commands
- Tools Perform Commands
- Tools ADP Commands

Tools Dump Commands

```
tools
    — dump
        — auto-discovery [detail] [log]
        — ppp port-id
        — router router-instance
            — 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-address |: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
            — ospf
                — 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]
            — rsvp
                — neighbor [ip-address] [detail]
                — psb [endpoint endpoint-address] [sender sender-address] [tunnelid
                  tunnel-id] [lspid lsp-id]
                — rsb [endpoint endpoint-address] [sender sender-address] [tunnelid
                  tunnel-id] [lspid lsp-id]
            — static-route
                — ldp-sync-status
        — 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]
        — ima
            — reset bundle-id
        — log
            — test-event
        — router router-instance
            — isis
                — ldp-sync-exit
                — run-manual-spf
            — mpls
                — cspf to ip-addr [from ip-addr] [bandwidth bandwidth] [include-bitmap bitmap] [exclude-bitmap bitmap] [hop-limit limit] [exclude-address excl-addr...(up to 8 max)] [use-te-metric] [strict-srlg] [srlg-group grp-id...(up to 8 max)] [exclude-node excl-node-id...(up to 8 max)] [skip-interface interface-name] [ds-class-type class-type] [cspf-reqtype req-type]
                — resignal {lsp lsp-name path path-name | delay minutes}
                — trap-suppress number-of-traps time-interval
            — ospf
                — ldp-sync-exit
                — 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
            — id service-id
                — endpoint endpoint-name
                    — force-switchover sdp-id:vc-id
                    — no force-switchover
```

Tools ADP Commands

tools
- — **auto-discovery** [**retry**] [**terminate**]
- — [**no**] **auto-discovery echo** [**debugger**]

Command Descriptions

Tools Dump Commands

Generic Commands

tools

Syntax **tools**

Context <root>

Description This command creates the context to enable useful tools for debugging purposes.

Default **n/a**

Dump Commands

dump

Syntax **dump**

Context tools

Description This command creates the context to display information for debugging purposes.

Default **n/a**

auto-discovery

Syntax **auto-discovery** [**detail**] [**log**]

Context tools>dump

Description This command allows you to view all progress and event logs stored by ADP.

Default **n/a**

Parameters **detail —** displays detailed information about the system, ports, and ADP instructions

log — displays all detailed progress and event logs with timestamps

ppp

Syntax **ppp** *port-id*

Context tools>dump

Description This command displays PPP information for a port.

Default **n/a**

Parameters *port-id* — specifies the port ID

Syntax: *port-id* *slot/mda/port*[.*channel*]
bundle bundle-*type-slot/mda.bundle-num*
bundle keyword
type ima, ppp
bundle-num 1 to 10

system-resources

Syntax **system-resources** *slot-number*

Context tools>dump

Description This command displays system resource information.

Default **n/a**

Parameters *slot-number* — specifies a specific slot to view system resources information

Dump Router Commands

router

Syntax **router** *router-instance*

Context tools>dump

Description This command enables tools for the router instance.

Default **n/a**

Parameters *router-instance* — specifies the router name and service ID

Values

router-name:	Base, management
service-id:	1 to 2147483647

Default Base

ldp

Syntax **ldp**

Context tools>dump>router

Description This command enables dump tools for LDP.

Default **n/a**

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 **n/a**

Parameters *ip-prefix/mask* — specifies the IP prefix and subnet mask

Values

ip-prefix:	a.b.c.d (host bits must be 0)
mask:	0 to 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 that 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*.

Values **Ethernet** — 0x0005

VLAN — 0x0004

vc-id — specifies the virtual circuit identifier

Values 1 to 4294967295

instance

Syntax **instance**

Context tools>dump>router>ldp

Description This command displays information for an LDP instance.

interface

Syntax **interface** [*ip-int-name* | *ip-address*]

Context tools>dump>router>ldp

Description This command displays information for an LDP interface.

Default **n/a**

Parameters *ip-int-name* — specifies the interface name

ip-address — specifies the IP address

memory-usage

Syntax **memory-usage**

Context tools>dump>router>ldp

Description This command displays memory usage information for LDP.

Default **n/a**

peer

Syntax **peer** *ip-address*

Context tools>dump>router>ldp

Description This command displays information for an LDP peer.

Default **n/a**

Parameters *ip-address* — specifies the IP address

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 **n/a**

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 **n/a**

timers

Syntax **timers**

Context tools>dump>router>ldp

Description This command displays timer information for LDP.

Default **n/a**

mpls

Syntax **mpls**

Context tools>dump>router

Description This command enables the context to display MPLS information.

Default **n/a**

ftn

Syntax **ftn** [**endpoint** *endpoint* | **sender** *sender* | **nexthop** *nexthop* | **lsp-id** *lsp-id* | **tunnel-id** *tunnel-id* | **label** *start-label end-label*]

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 **n/a**

Parameters *endpoint* — specifies the IP address of the last hop

Values a.b.c.d

sender — specifies the IP address of the sender

Values a.b.c.d

nexthop — specifies the IP address of the next hop

Values a.b.c.d

lsp-id — specifies the label switched path that is signaled for this entry

Values 0 to 65535

tunnel-id — specifies the SDP ID

Values 0 to 65535

start-label end-label — specifies the label range for the information dump

Values start-label — 32 to 131071

end-label — 32 to 131071

ilm

Syntax **ilm** [**endpoint** *endpoint* | **sender** *sender* | **nexthop** *nexthop* | **lsp-id** *lsp-id* | **tunnel-id** *tunnel-id* | **label** *start-label end-label*]

Context tools>dump>router>mpls

Description This command displays incoming label map (ILM) information for MPLS.

Default **n/a**

Parameters *endpoint* — specifies the IP address of the last hop

Values a.b.c.d

sender — specifies the IP address of the sender

Values a.b.c.d

nexthop — specifies the IP address of the next hop

Values a.b.c.d

lsp-id — specifies the label switched path that is signaled for this entry

Values 0 to 65535

tunnel-id — specifies the SDP ID

Values 0 to 65535

start-label end-label — specifies the label range for the information dump

Values start-label — 32 to 131071

end-label — 32 to 131071

lspinfo

Syntax **lspinfo** [*lsp-name*] [**detail**]

Context tools>dump>router>mpls

Description This command displays LSP information for MPLS.

Default **n/a**

Parameters *lsp-name* — the LSP identifier

Values up to 32 characters (must be unique)

detail — displays detailed LSP information

memory-usage

Syntax **memory-usage**

Context tools>dump>router>mpls

Description This command displays memory usage information for MPLS.

Default **n/a**

ospf

Syntax **ospf**

Context tools>dump>router

Description This command enables the context to display tools information for OSPF.

Default **n/a**

abr

Syntax **abr** [**detail**]

Context tools>dump>router>ospf

Description This command displays area border router (ABR) information for OSPF.

Default **n/a**

Parameters **detail —** displays detailed information about the ABR

asbr

Syntax **asbr** [**detail**]

Context tools>dump>router>ospf

Description This command displays autonomous system boundary router (ASBR) information for OSPF.

Default **n/a**

Parameters **detail —** displays detailed information about the ASBR

bad-packet

Syntax **bad-packet** [*interface-name*]

Context tools>dump>router>ospf

Description This command displays information about bad packets for OSPF.

Default **n/a**

Parameters *interface-name* — displays only the bad packets identified by this interface name

leaked-routes

Syntax **leaked-routes** [**summary** | **detail**]

Context tools>dump>router>ospf

Description This command displays information about leaked routes for OSPF.

Default **summary**

Parameters **summary —** displays a summary of information about leaked routes for OSPF

detail — displays detailed information about leaked routes for OSPF

memory-usage

Syntax **memory-usage** [**detail**]

Context tools>dump>router>ospf

Description This command displays memory usage information for OSPF.

Default **n/a**

Parameters **detail —** displays detailed information about memory usage for OSPF

request-list

Syntax **request-list** [**neighbor** *ip-address*] [**detail**]
request-list [**virtual-neighbor** *ip-address* **area-id** *area-id]* [**detail**]

Context tools>dump>router>ospf

Description This command displays request list information for OSPF.

Default **n/a**

Parameters **neighbor** *ip-address* — displays neighbor information only for the neighbor identified by the IP address

detail — displays detailed information about the neighbor or virtual neighbor

virtual-neighbor *ip-address* — displays information about the virtual neighbor identified by the IP address

area-id — the OSPF area ID expressed in dotted-decimal notation or as a 32-bit decimal integer

retransmission-list

Syntax **retransmission-list** [**neighbor** *ip-address*] [**detail**]
retransmission-list [**virtual-neighbor** *ip-address* **area-id** *area-id*] [**detail**]

Context tools>dump>router>ospf

Description This command displays dump retransmission list information for OSPF.

Default **n/a**

Parameters **neighbor** *ip-address* — displays neighbor information only for the neighbor identified by the IP address

detail — displays detailed information about the neighbor or virtual neighbor

virtual-neighbor *ip-address* — displays information about the virtual neighbor identified by the IP address

area-id — the OSPF area ID expressed in dotted-decimal notation or as a 32-bit decimal integer

route-summary

Syntax **route-summary**

Context tools>dump>router>ospf

Description This command displays dump route summary information for OSPF.

Default **n/a**

route-table

Syntax **route-table** [**type**] [**detail**]

Context tools>dump>router>ospf

Description This command displays dump information about routes learned through OSPF.

Default **n/a**

Parameters **type —** the type of route table to display information about

Values intra-area, inter-area, external-1, external-2, nssa-1, nssa-2

detail — displays detailed information about learned routes

rsvp

Syntax **rsvp**

Context tools>dump>router

Description This command enables the context to display tools information for RSVP.

Default **n/a**

neighbor

Syntax **neighbor** [*ip-address*] [**detail**]

Context tools>dump>router>rsvp

Description This command displays neighbor information for RSVP.

Default **n/a**

Parameters *ip-address* — the IP address of the neighbor

Values a.b.c.d

detail — displays detailed information about the neighbor

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 **n/a**

Parameters *endpoint-address* — specifies the IP address of the last hop

sender-address — specifies the IP address of the sender

tunnel-id — specifies the SDP ID

Values 0 to 4294967295

lsp-id — specifies the label switched path that is signaled for this entry

Values 1 to 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 **n/a**

Parameters *endpoint-address* — specifies the IP address of the last hop

sender-address — specifies the IP address of the sender

tunnel-id — specifies the SDP ID

Values 0 to 4294967295

lsp-id — specifies the label switched path that is signaled for this entry

Values 1 to 65535

static-route

Syntax **static-route**

Context tools>dump>router

Description This command enables the context to display tools information for static routes.

Default **n/a**

ldp-sync-status

Syntax **ldp-sync-status**

Context tools>dump>router>static-route

Description This command displays the status of the LDP synchronization timers for static routes.

Tools Perform Commands

Perform Commands

perform

Syntax **perform**

Context tools

Description This command enables the context to specify tools to perform specific tasks.

Default **n/a**

aps

Syntax **aps**

Context tools>perform

Description This command enables the context to perform APS operations.

clear

Syntax **clear aps**-*id* {**protect** | **working**}

Context tools>perform>aps

Description This command removes all APS operational commands.

Parameters **aps**-*id* — the specified APS group

protect — the physical port acting as a protection circuit for the APS group

working — the physical port acting as a working circuit for the APS group

exercise

Syntax **exercise aps**-*id* {**protect** | **working**}

Context tools>perform>aps

Description This command performs an exercise request on the protection or working circuit.

Parameters **aps**-*id* — the specified APS group

protect — the physical port acting as a protection circuit for the APS group

working — the physical port acting as a working circuit for the APS group

force

Syntax **force aps**-*id* {**protect** | **working**}

Context tools>perform>aps

Description This command forces a switch to either the protection or working circuit.

Parameters **aps**-*id* — the specified APS group

protect — the physical port acting as a protection circuit for the APS group

working — the physical port acting as a working circuit for the APS group

lockout

Syntax **lockout aps**-*id*

Context tools>perform>aps

Description This command locks out the protection circuit in the specified APS group.

Parameters **aps**-*id* — the specified APS group

request

Syntax **request aps**-*id* {**protect** | **working**}

Context tools>perform>aps

Description This command requests a manual switch to either the protection or working circuit.

Parameters **aps**-*id* — the specified APS group

protect — the physical port acting as a protection circuit for the APS group

working — the physical port acting as a working circuit for the APS group

cron

Syntax **cron**

Context tools>perform

Description This command enables the context to perform CRON (scheduling) control operations.

Default **n/a**

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

action-owner — specifies the owner name

Default TiMOS CLI

all — specifies to stop all CRON scripts

ima

Syntax **ima**

Context tools>perform

Description This command enables the context to perform IMA operations.

Default **n/a**

reset

Syntax **reset** *bundle-id*

Context tools>perform>ima

Description This command resets an IMA bundle in the startup state.

Default **n/a**

Parameters *bundle-id* — specifies the IMA bundle ID

Syntax: bundle-ima-*slot/mda.bundle-num*
bundle-ima keyword
bundle-num 1 to 10

log

Syntax **log**

Context tools>perform

Description This command enables event logging tools.

test-event

Syntax **test-event**

Context tools>perform>log

Description This command generates a test event.

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 **server-address** *ip-address* — specifies the server ID

Values a.b.c.d

port — specifies the port ID

Values 1 to 65535

dhcp-client-user-name — specifies the DHCP client

Values 256 characters maximum

password — specifies the CLI access password

Values 10 characters maximum

key — specifies the authenication key

Values 20 characters maximum

source-address *ip-address* **—** specifies the source IP address of the DHCP relay messages

Values a.b.c.d

seconds — specifies the timeout in seconds

Values 1 to 90

router-instance — specifies the router name or service ID

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

Default Base

service

Syntax **service**

Context tools>perform

Description This command enables the context to configure tools for services.

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* — specifies an existing service ID

Values 1 to 2147483647

endpoint

Syntax **endpoint** *endpoint-name*

Context tools>perform>service>id

Description This command enables the context to configure tools for a specific service endpoint.

Parameters *endpoint-name* — specifies an existing 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* — specifies an existing spoke SDP for the service

Values

sdp-id:	1 to 17407
vc-id:	1 to 4294967295

Perform Router Commands

router

Syntax **router** *router-instance*

Context tools>perform

Description This command enables tools for the router instance.

Default **n/a**

Parameters *router-instance* — specifies the router name and service ID

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

Default Base

isis

Syntax **isis**

Context tools>perform>router

Description This command enables the context to perform specific IS-IS tasks.

mpls

Syntax **mpls**

Context tools>perform>router

Description This command enables the context to perform specific MPLS tasks.

Default **n/a**

cspf

Syntax **cspf to** *ip-addr* [**from** *ip-addr*] [**bandwidth** *bandwidth*] [**include-bitmap** *bitmap*] [**exclude-bitmap** *bitmap*] [**hop-limit** *limit*] [**exclude-address** *excl-addr*...(up to 8 max)] [**use-te-metric**] [**strict-srlg**] [**srlg-group** *grp-id*...(up to 8 max)] [**exclude-node** *excl-node-id*...(up to 8 max)] [**skip-interface** *interface-name*] [**ds-class-type** *class-type*] [**cspf-req-type** *req-type*]

Context tools>perform>router>mpls

Description This command computes a CSPF path with specified user constraints.

Default **n/a**

Parameters **to** *ip-addr* **—** the destination IP address

Values a.b.c.d

from *ip-addr* **—** the originating IP address

Values a.b.c.d

bandwidth — the amount of bandwidth in megabits per second (Mb/s) to be reserved

Values 0 to 4294967295 (values can be expressed in decimal, hexadecimal, or binary)

include-bitmap *bitmap* **—** specifies to include a bitmap that lists the admin groups that should be included during the CSPF computation

exclude-bitmap *bitmap* **—** specifies to exclude a bitmap that lists the admin groups that should be included during the CSPF computation

limit — the total number of hops an FRR bypass LSP can take before merging back onto the main LSP path

Values 1 to 255

excl-addr — an IP address to exclude from the CSPF computation (up to a maximum of eight addresses in one command)

Values a.b.c.d (outbound interface)

use-te-metric — specifies to use the traffic engineering metric used on the interface

strict-srlg — specifies to use strict **frr-srlg** to compute a new CSPF path

grp-id — specifies to use up to eight SRLGs to compute a new CSPF path

Values 0 to 4294967295

excl-node-id — a node to exclude from the CSPF computation (up to a maximum of eight nodes in one command)

Values a.b.c.d

interface-name — a local interface name (rather than the address) to exclude from the CSPF computation

Values max 32 characters

class-type — the class type

Values 0 to 7

req-type — the CSPF request type

Values all – all ECMP paths

random – random ECMP paths

resignal

Syntax **resignal** {**lsp** *lsp-name* **path** *path-name* | **delay** *minutes*}

Context tools>perform>router>mpls

Description This command resignals specified LSP paths. The *minutes* parameter is used to configure the global timer to resignal all LSPs. The resignal timer is the time before resignaling occurs after the resignal condition occurs. If only *lsp-name* and *path-name* are provided, the specified LSP is resignaled immediately. For the delay option to work, the resignal time in the **configure>router>mpls** context must be set.

Default **n/a**

Parameters *lsp-name* — specifies a unique LSP name, up to 32 characters

path-name — specifies the name for the LSP path, up to 32 characters

minutes — specifies the delay interval, in minutes, before all LSPs are resignaled. If the value 0 is entered, all LSPs are resignaled immediately.

Values 0 to 30

trap-suppress

Syntax **trap-suppress** *number-of-traps time-interval*

Context tools>perform>router>mpls

Description This command modifies thresholds for trap suppression. The command is used to suppress traps after the specified number of traps has been raised within the specified period of time.

Default **n/a**

Parameters *number-of-traps* — specifies the number of traps in multiples of 100. An error message is generated if an invalid value is entered.

Values 100 to 1000

time-interval — specifies the time interval in seconds

Values 1 to 300

ospf

Syntax **ospf**

Context tools>perform>router

Description This command enables the context to perform specific OSPF tasks.

ldp-sync-exit

Syntax **ldp-sync-exit**

Context tools>perform>router>ospf
tools>perform>router>isis

Description This command terminates IGP-LDP synchronization. OSPF or IS-IS then advertises the actual cost value of the link for all interfaces that have IGP-LDP synchronization enabled, if the currently advertised cost is different.

refresh-lsas

Syntax **refresh-lsas** [*lsa-type*] [*area-id*]

Context tools>perform>router>ospf

Description This command refreshes LSAs for OSPF.

Parameters *lsa-type* — the specified LSA type

Values router, network, summary, asbr, extern, nssa, opaque

area-id — the OSPF area ID expressed in dotted-decimal notation or as a 32-bit integer

Values 0.0.0.0 to 255.255.255.255 (dotted-decimal), 0 to 4294967295 (decimal integer)

run-manual-spf

Syntax **run-manual-spf** [**externals-only**]

Context tools>perform>router>ospf
tools>perform>router>isis

Description This command runs the shortest path first (SPF) algorithm for OSPF or IS-IS.

The **externals-only** parameter applies only to OSPF.

Parameters **externals-only —** specifies the route preference for OSPF external routes

Tools ADP Commands

auto-discovery

Syntax **auto-discovery** [**retry**] [**terminate**]

Context tools

Description This command is used to control ADP while it is running.

The **retry** keyword restarts ADP if it has been halted due to errors. Executing this command clears the rejected DHCP server list for all ports and retries any processing that failed.

The **terminate** keyword terminates ADP and removes the ADP keyword from the BOF. The router returns to normal operations and any temporary configuration is removed. Network configuration and remote access remain enabled to allow the router to be manually provisioned remotely. ADP will not run again on future system restarts unless it is re-enabled via the CLI.

Default **n/a**

Parameters **retry —** resumes ADP after being halted for errors

terminate — terminates ADP and removes the ADP keyword from the BOF

auto-discovery echo

Syntax [**no**] **auto-discovery echo** [**debugger**]

Context tools

Description This command enables ADP echoing, which sends periodic updates to the console. The default is for ADP to echo progress summaries and major events. For troubleshooting, the optional **debugger** parameter causes ADP to echo detailed progress reports with events and timestamps. The command reverts to the default settings each time ADP is run on the system.

The **no** form of this command disables ADP echoing.

Default **auto-discovery echo**

Parameters **debugger —** enables ADP echoing of detailed progress reports with events and timestamps

Standards and Protocol Support

Standards Compliance

IEEE 802.1ag Service Layer OAM
IEEE 802.1p/q VLAN Tagging
IEEE 802.3 10BaseT
IEEE 802.3ah Ethernet OAM
IEEE 802.3u 100BaseTX
IEEE 802.3x Flow Control
IEEE 802.3z 1000BaseSX/LX
IEEE 802.3-2008 Revised base standard
ITU-T Y.1731 OAM functions and mechanisms for Ethernet-based networks

Telecom Compliance

IC CS-03 Issue 9 Spectrum Management and Telecommunications
ACTA TIA-968-A
AS/ACIF S016 (Australia/New Zealand) Requirements for Customer Equipment for connection to hierarchical digital interfaces
ITU-T G.703 Physical/electrical characteristics of hierarchical digital interfaces
ITU-T G.707 Network node interface for the Synchronous Digital Hierarchy (SDH)
ITU-T G.712-2001 Transmission performance characteristics of pulse code modulation channels
ITU-T G.957 Optical interfaces for equipments and systems relating to the synchronous digital hierarchy
ITU-T V.24 List of definitions for interchange circuits between data terminal equipment (DTE) and data circuit-terminating equipment (DCE)
ITU-T V.36 Modems for synchronous data transmission using 60-108 kHz group band circuits
ITU-T X.21 Interface between Data Terminal Equipment and Data Circuit-Terminating Equipment for Synchronous Operation on Public Data Networks

Protocol Support

ATM

RFC 2514 Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management, February 1999
RFC 2515 Definition of Managed Objects for ATM Management, February 1999
RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5
af-tm-0121.000 Traffic Management Specification Version 4.1, March 1999
ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
ITU-T Recommendation I.432.1 - B-ISDN user-network interface - Physical layer specification: General characteristics
GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994
AF-PHY-0086.001 Inverse Multiplexing for ATM (IMA)

BFD

draft-ietf-bfd-mib-00.txt Bidirectional Forwarding Detection Management Information Base
draft-ietf-bfd-base-o5.txt Bidirectional Forwarding Detection
draft-ietf-bfd-v4v6-1hop-06.txt BFD IPv4 and IPv6 (Single Hop)
draft-ietf-bfd-multihop-06.txt BFD for Multi-hop Paths

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MDS
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 BGP-4
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs) (previously RFC 2574bis BGP/MPLS VPNs)
RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP (previously RFC 1966 and RFC 2796)
RFC 4724 Graceful Restart Mechanism for BGP - GR Helper
RFC 4760 Multi-protocol Extensions for BGP (previously RFC 2858)
RFC 4893 BGP Support for Four-octet AS Number Space

DHCP/DHCPv6

RFC 1534 Interoperation between DHCP and BOOTP
RFC 2131 Dynamic Host Configuration Protocol (REV)
RFC 3046 DHCP Relay Agent Information Option (Option 82)
RFC 3315 Dynamic Host Configuration Protocol for IPv6

DIFFERENTIATED SERVICES

RFC 2474 Definition of the DS Field in the IPv4 and IPv6 Headers
RFC 2597 Assured Forwarding PHB Group
RFC 2598 An Expedited Forwarding PHB
RFC 3140 Per-Hop Behavior Identification Codes

DIGITAL DATA NETWORK MANAGEMENT

V.35
RS-232 (also known as EIA/TIA-232)

GRE

RFC 2784 Generic Routing Encapsulation (GRE)

IPv6

RFC 2460 Internet Protocol, Version 6 (IPv6) Specification
RFC 2462 IPv6 Stateless Address Autoconfiguration
RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
RFC 3587 IPv6 Global Unicast Address Format
RFC 3595 Textual Conventions for IPv6 Flow Label
RFC 4007 IPv6 Scoped Address Architecture
RFC 4193 Unique Local IPv6 Unicast Addresses
RFC 4291 IPv6 Addressing Architecture
RFC 4443 Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification
RFC 4649 DHCPv6 Relay Agent Remote-ID Option
RFC 4861 Neighbor Discovery for IP version 6 (IPv6)

LDP

RFC 5036 LDP Specification

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 (IS-IS) 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 4205 for Shared Risk Link Group (SRLG) TLV draft-ietf-isis-igp-p2p-over-lan-05.txt
RFC 5309 Point-to-Point Operation over LAN in Link State Routing Protocols

MPLS

RFC 3031 MPLS Architecture
RFC 3032 MPLS Label Stack Encoding
RFC 3815 Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)
RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

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 1305 Network Time Protocol (Version 3) Specification, Implementation and Analysis
RFC 1850 OSPF-MIB
RFC 1907 SNMPv2-MIB
RFC 2011 IP-MIB
RFC 2012 TCP-MIB
RFC 2013 UDP-MIB
RFC 2030 Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI
RFC 2096 IP-FORWARD-MIB
RFC 2138 RADIUS
RFC 2206 RSVP-MIB
RFC 2571 SNMP-FRAMEWORKMIB
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 2588 SONET-MIB
RFC 2665 EtherLike-MIB
RFC 2819 RMON-MIB
RFC 2863 IF-MIB
RFC 2864 INVERTED-STACK-MIB
RFC 3014 NOTIFICATION-LOG MIB
RFC 3164 The BSD Syslog Protocol
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
draft-ietf-disman-alarm-mib-04.txt
draft-ietf-mpls-ldp-mib-07.txt
draft-ietf-ospf-mib-update-04.txt
draft-ietf-mpls-lsr-mib-06.txt
draft-ietf-mpls-te-mib-04.txt
IANA-IFType-MIB

OSPF

RFC 1765 OSPF Database Overflow
RFC 2328 OSPF Version 2
RFC 2370 Opaque LSA Support
RFC 3101 OSPF NSSA Option
RFC 3137 OSPF Stub Router Advertisement
RFC 3630 Traffic Engineering (TE) Extensions to OSPF
RFC 4203 Shared Risk Link Group (SRLG) sub-TLV

PPP

RFC 1332 PPP Internet Protocol Control Protocol (IPCP)
RFC 1570 PPP LCP Extensions
RFC 1619 PPP over SONET/SDH
RFC 1661 The Point-to-Point Protocol (PPP)
RFC 1662 PPP in HDLC-like Framing
RFC 1989 PPP Link Quality Monitoring
RFC 1990 The PPP Multilink Protocol (MP)
RFC 2686 The Multi-Class Extension to Multi-Link PPP

PSEUDOWIRES

RFC 3550 RTP: A Transport Protocol for Real-Time Applications
RFC 3985 Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture
RFC 4385 Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN
RFC 4446 IANA Allocation for PWE3
RFC 4447 Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)

RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks
RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)
RFC 4717 Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks
RFC 5085 Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires
RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)
draft-ietf-pwe3-redundancy-02 Pseudowire (PW) Redundancy

RADIUS

RFC 2865 Remote Authentication Dial In User Service
RFC 2866 RADIUS Accounting

RSVP-TE and FRR

RFC 2430 A Provider Architecture for DiffServ & TE
RFC 2961 RSVP Refresh Overhead Reduction Extensions
RFC 2702 Requirements for Traffic Engineering over MPLS
RFC 2747 RSVP Cryptographic Authentication
RFC 3097 RSVP Cryptographic Authentication - Updated Message Type Value
RFC 3209 Extensions to RSVP for LSP Tunnels
RFC 3210 Applicability Statement for Extensions to RSVP for LSP Tunnels
RFC 4090 Fast Reroute Extensions to RSVP-TE for LSP Tunnels

SONET/SDH

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000
ITU-T Recommendation 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

SSH

draft-ietf-secsh-architecture.txt SSH 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

SYNCHRONIZATION

G.813 Timing characteristics of SDH equipment slave clocks (SEC)
G.8261 Timing and synchronization aspects in packet networks
G.8262 Timing characteristics of synchronous Ethernet equipment slave clock
GR 1244 CORE Clocks for the Synchronized Network: Common Generic Criteria
IEEE 1588v2 1588 PTP 2008

TACACS+

IETF draft-grant-tacacs-02.txt The TACACS+ Protocol

TCP/IP

RFC 768 User Datagram Protocol
RFC 791 Internet Protocol
RFC 792 Internet Control Message Protocol
RFC 793 Transmission Control Protocol
RFC 826 Ethernet Address Resolution Protocol
RFC 854 Telnet Protocol Specification
RFC 1350 The TFTP Protocol (Rev. 2)
RFC 1812 Requirements for IPv4 Routers

VPLS

RFC 4762 Virtual Private LAN Services Using LDP

Proprietary MIBs

TIMETRA-ATM-MIB.mib
TIMETRA-CAPABILITY-7705-V1.mib
TIMETRA-CFLOWD-MIB.mib
TIMETRA-CHASSIS-MIB.mib
TIMETRA-CLEAR-MIB.mib
TIMETRA-FILTER-MIB.mib
TIMETRA-GLOBAL-MIB.mib
TIMETRA-LDP-MIB.mib
TIMETRA-LOG-MIB.mib
TIMETRA-MPLS-MIB.mib
TIMETRA-OAM-TEST-MIB.mib
TIMETRA-PORT-MIB.mib
TIMETRA-PPP-MIB.mib
TIMETRA-QOS-MIB.mib
TIMETRA-ROUTE-POLICY-MIB.mib
TIMETRA-RSVP-MIB.mib
TIMETRA-SAP-MIB.mib
TIMETRA-SDP-MIB.mib
TIMETRA-SECURITY-MIB.mib
TIMETRA-SERV-MIB.mib
TIMETRA-SYSTEM-MIB.mib
TIMETRA-TC-MIB.mib

Customer documentation and product support

Customer documentation

http://www.alcatel-lucent.com/myaccess

Product manuals and documentation updates are available at alcatel-lucent.com. If you are a new user and require access to this service, please contact your Alcatel-Lucent sales representative.

Technical Support

http://support.alcatel-lucent.com

Documentation feedback

documentation.feedback@alcatel-lucent.com

3HE 05975 AAAA TQZZA Edition 01