JDSU

White Paper

Ethernet OAM Test Applications*

By Reza Vaez-Ghaemi, Ph.D.

The rising level of Triple Play, Internet Protocol (IP) Business, and 3G wireless service deployments drives the demand for higher bandwidth and cost efficiency. Carrier Ethernet promises to deliver many attributes necessary for a cost-effective, carrier-grade technology in telecom networks. For Ethernet to become carrier-grade, a number of improvements must be developed and deployed. This paper introduces operation, administration, and maintenance (OAM), and examines some test applications for its qualification before deployment in field.

Introduction

Competition for market share among telecom, cable/multimedia service operator (MSO), and mobile operators further fuels the need for network investments. When compared to voice-only service delivery, average revenue per user is expected to rise drastically with IP television (IPTV) and higher-speed data services. Telecom operators are investing heavily in Triple-Play services, particularly IPTV and related access networks capable of carrying increased data rates of typically 40-50 Mb/s minimum, which will grow over time as newer high-bandwidth services are offered, such as three-dimensional (3D) video.

The Ethernet community has been continuously adding capabilities to native Ethernet in order to deliver a carrier-grade transport alternative to legacy time division multiplexed (TDM) technologies. Provider Backbone Bridge/Provider Backbone Transport (PBB/PBT), Transport Multi-Protocol Label Switching/Multi-Protocol Label Switching Transport Profile (T-MPLS/MPLS-TP), Ethernet OAM, Synchronous Ethernet, and Time Division Multiplexing over IP (TDMoIP) are some examples of major initiatives for carrier-grade Ethernet. These technologies require new methods and tools for qualification in labs and in the field.

Attributes of Carrier Ethernet

Carrier Ethernet networks offer transport for millions of end users and connectivity among final users as well as among client networks. Native Ethernet was designed with the target of servicing much smaller groups of users in campuses and enterprises, but carrier-grade Ethernet (Figure 1) must deliver scalability similar to conventional carrier networks.

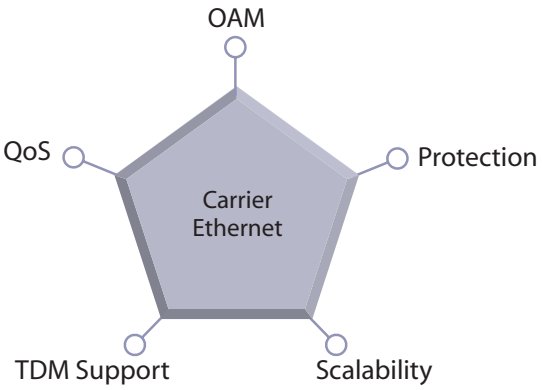

Figure 1: Requirements for Carrier Ethernet

The simultaneous use of services by large numbers of subscribers and applications leads to resource conflicts and congestion. Some applications are more critical and, therefore, are delivered at premium prices. Carrier Ethernet must provide service differentiation, so that those critical applications receive the quality necessary.

Another attribute related to these critical applications revolves around protection and restoration. In the event of hard failures, users expect protection schemes that can deliver performance at or below 50 ms. Also, in the event of failures, rapid fault identification and localization is key in large networks driving the demand for carrier-grade OAM functions.

Finally, while packet traffic represents the largest growing segment in user traffic, the demand for TDM traffic remains robust. Thus, any carrier-grade technology must transport TDM-based traffic at performance levels similar to circuit-based technologies.

Ethernet OAM

Ethernet OAM provides a key challenge before Ethernet is mass marketed in carrier networks; however, standards exist to address most of the issues. The new standards will provide new OAM capability to the customer premises demarcation point, potentially reducing operating expenses (OPEX) by more than half. Low Ethernet capital expenditures (CAPEX) shift the profitability focus to OPEX, and Ethernet OAM is key to managing Ethernet service OPEX.

Ethernet OAM functions are provided in various layers and segments of the network, as Figure 2 shows:

– Access layer
– End-to-end connectivity fault management (CFM)
– Service layer

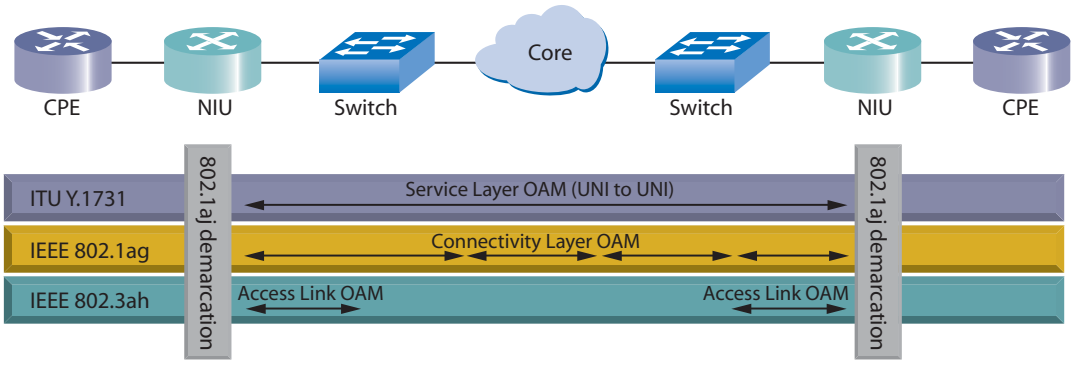

Figure 2: Overview Ethernet OAM

Figure 3 shows that multiple standards provide building blocks for managing the access:

– 802.3ah OAM addresses managing the physical layer from the provider edge (PE) to the media access control (MAC) layer of the remote device
– Metro Ethernet Forum (MEF) Ethernet Local Management Interface (E-LMI) addresses managing the User Network Interface (UNI) of the remote device
– 802.1aj Two-Port MAC Relay (TPMR) addresses managing a customer demarcation device

The IEEE 802.1aj TPMR defines protocols for interactions between provider bridges and remote two-port relay devices that might be used for demarcation. It uses simple network management protocol (SNMP) natively over Ethernet without using IP. TPMR enables configuration of remote parameters such as virtual local area network (VLAN), class of service (CoS), and quality of service (QoS). It also delivers status information, if the demarcation device is a two-port relay forwarder.

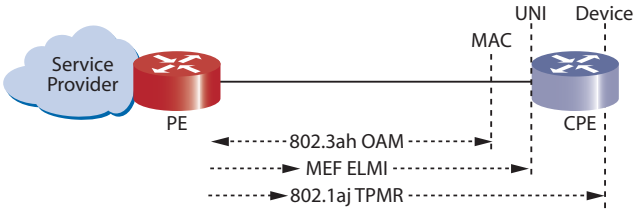

Figure 3: Managing Edge Devices

Ethernet First Mile (802.3ah) Link Layer OAM

The Ethernet First Mile (EFM) 802.3ah provides link layer OAM functionality in the first/last mile, as Figure 4 demonstrates. It is media independent and operates at a slow rate of 10 frames per second. Ethernet OAM packet data units (OAMPDU) only work in point-to-point full-duplex networks and are not forwarded by peer devices. They require minimal configuration and deliver following functions:

– Device discovery

– Remote failure indication

– Remote loopback

– Link monitoring

During the network initialization, adjacent devices exchange identification information and OAM capabilities. With remote failure indication, network devices can notify peer devices in the event of failures. The remote loopback is a link-layer mechanism that operates at the frame level. Link monitoring delivers event notifications, such as status and diagnostics information, that are stored in local management information bases (MIB), where peers can pull them.

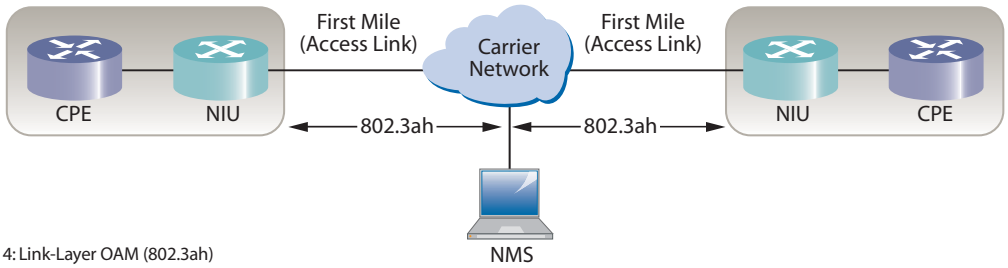

Figure 4: Link-Layer OAM (802.3ah)

OAM provides an optional data link layer frame-level loopback mode, which is controlled remotely. OAM remote loopback can be used for fault localization and link performance testing. Statistics from both the local and remote data terminal equipment (DTE) can be queried and compared at any time while the remote DTE is in OAM remote loopback mode. These queries can take place before, during, or after loopback frames have been sent to the remote DTE. In addition, an OAM implementation may analyze loopback frames within the OAM sublayer to determine additional information about the health of the link, for instance determining which frames are being dropped due to link errors). Figure 5 shows the path of frames traversing the layer stack of both the local and remote DTE.

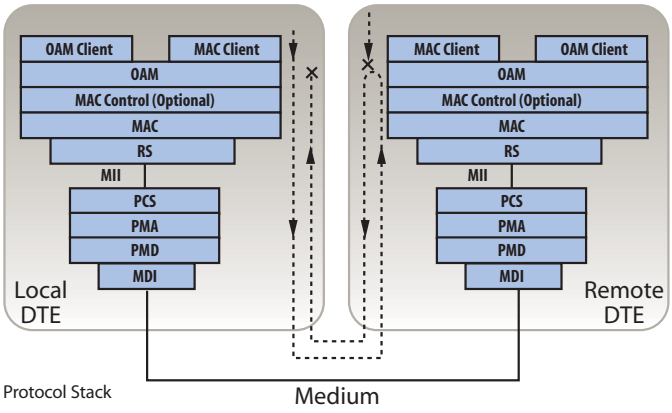

Figure 5: Ethernet OAM in Protocol Stack

Connectivity Fault Management (802.1ag)*

The Connectivity Fault Management (CFM) standard specifies protocols and protocol entities within the architecture of VLAN-aware bridges that enable the detection, verification, and isolation of connectivity failures in virtual bridged LANs (VBLAN). These capabilities can be used in networks operated by multiple independent organizations, each with restricted management access to each other's equipment.

This standard specifies protocols, procedures, and managed objects in support of connectivity fault management. These allow discovery and verification of the path through bridges and LANs taken from frames addressed to and from specified network users. It enables detection and isolation of a connectivity fault to a specific bridge or LAN. The standard:

– Defines maintenance domains, maintenance associations, their constituent maintenance points, and the managed objects required to create and administer them.

– Describes the protocols and procedures that maintenance points use to detect and diagnose connectivity faults within a maintenance domain (MD).

Maintenance Domain (MD): The network or the part of the network for which faults in connectivity can be managed.

Connectivity Fault Management (CFM): Comprises capabilities for detecting, verifying, and isolating connectivity failures in VLANs.

Maintenance Entity (ME): Represents an entity that requires management and facilitates a relationship between two ME group end points.

ME Group (MEG): Includes different MEs that satisfy the following conditions:

– MEs in a MEG exist in the same administrative boundary

– MEs in a MEG have the same MEG level

– MEs in a MEG belong to the same point-to-point or multipoint Ethernet connections.

MEG End Point (MEP): Marks the end point of an Ethernet MEG that can initiate and terminate OAM frames for fault management and performance monitoring.

MEG Intermediate Point (MIP): Serves as an intermediate point in a MEG that reacts to certain OAM frames. A MIP does not initiate OAM frames, nor does it take action on the transit Ethernet flows.

Maintenance Association (MA): A set of MEPs that are each configured with the same maintenance association identifier (MAID) and MD level, which are established to verify the integrity of a single service instance.

Maintenance Association End-Point Identifier (MEPID): A small integer, unique over a given MA, which identifies a specific MEP.

Maintenance Association Identifier (MAID): An identifier for an MA, unique over the domain, that uses CFM to protect against the accidental concatenate.

OAM transparency allows for transparent carrying of OAM frames belonging to higher-level MEGs across other lower-level MEGs when the MEGs are nested.

OAM frames belonging to an administrative domain originate and terminate in MEPs present at the boundary of that administrative domain. A MEP prevents OAM frames that correspond to a MEG in the administrative domain from leaking outside of that administrative domain. However, when a MEP is not present or is faulty, the associated OAM frames can leave the administrative domain.

Similarly, a MEP present at the boundary of an administrative domain protects the administrative domain from OAM frames belonging to other administrative domains. The MEP allows OAM frames from outside administrative domains belonging to higher-level MEs to pass transparently, while it blocks OAM frames from outside administrative domains belonging to same or lower-level MEs.

CFM functions are partitioned as follows:

– Path discovery
– Fault detection
– Fault verification and isolation
– Fault notification
– Fault recovery

Path discovery uses the Linktrace Protocol to determine the path taken to a target MAC address. A Linktrace Message (LTM) is multicast from a MEP to its neighboring MIPs, and from MIP to MIP, to the multipoint processor (MP) terminating the path. Each MIP along the path and the terminating MP return unicast Linktrace Replies (LTR) to the originating MEP.

Fault detection uses the Continuity Check Protocol to detect both connectivity failures and unintended connectivity between service instances. Each MEP can periodically transmit a multicast Connectivity Check Message (CCM) that announces the identity of the MEP and its MA and tracks the CCMs received from the other MEPs. Connectivity faults that can misdirect a CCM are revealed as differences between the CCMs received and the configured expectations of the MEP.

Fault verification and fault isolation are administrative actions that are typically performed after fault detection. Fault verification also can confirm successful initiation or restoration of connectivity. The administrator uses the Loopback Protocol to perform fault verification.

The MEP provides fault notification when a connectivity fault is detected in its MA, either because expected CCMs were not received, unexpected or invalid CCMs were received, or a CCM carried a notification of the failure of its associated bridge port.

Fault recovery is provided by the Spanning Tree Protocols and by activities of the network administrator, such as the correction of configuration errors, or replacement of failed components that are outside the scope of this standard.

Eight MEG levels are available to accommodate different network deployment scenarios. When the data path flows of customers, providers, and operators are indistinguishable, based on means of the Ethernet layer encapsulations, the eight MEG levels can be shared among them to distinguish between OAM frames belonging to nested MEGs of each. The default MEG level assignment among the roles of customers, providers, and operators is:

– Customer role is assigned three MEG levels: 7, 6, and 5
– Provider role is assigned two MEG levels: 4 and 3
– Operator role is assigned three MEG levels: 2, 1, and 0

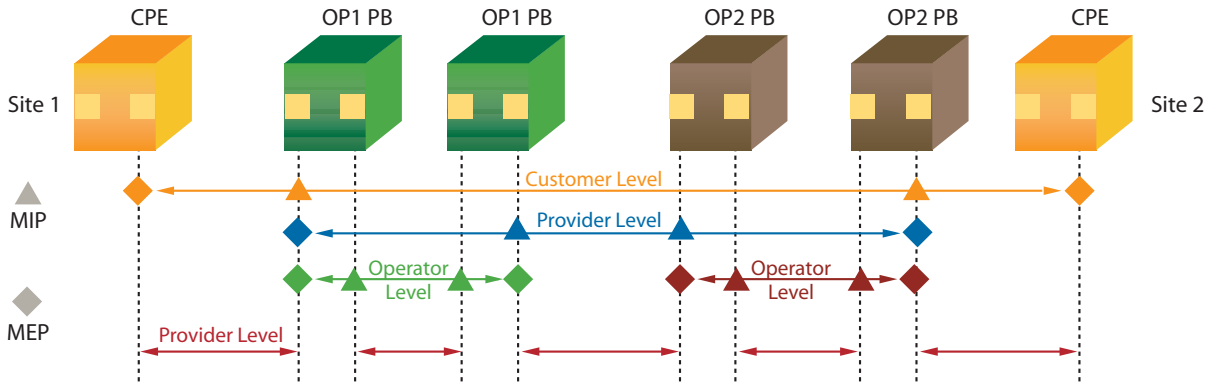

Figure 6: Maintenance Domains and Entities

Continuity Check

To verify CFM it is important to first properly set up these maintenance domains and associations:

– MD
– MA
– CCM interval rate
– MEPID for local and remote device

Figure 7 illustrates that CCM verification is among the most important CFM tests because it:

– Checks the CCM interval rate

– Compares the received CCM interval rate against the expected CCM interval rate

– Checks for missing messages

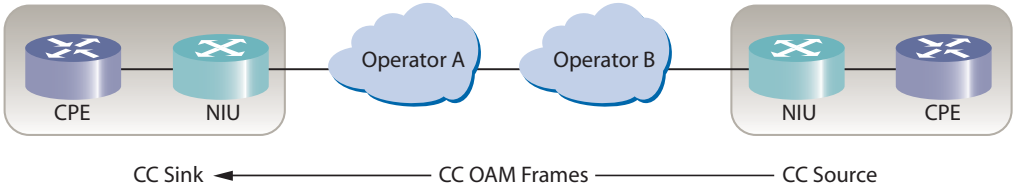

Figure 7: Continuity Check Messages

When emulating a MEP, the tester collects the CCM and, based on this message, can then calculate the CCM interval/rate. The tester collects only CCM equal to or less than the specified MD level with the same MEPID/MAID. A CCM tester will behave the same; for example, if the use case is about MD level 4, the tester collects CCM for MD levels 1 thru 4:

– If the tester receives a CCM with a MD level lower than the expected MD level, it detects an unexpected MA/MD level.

– If the tester receives a CCM with same MD level, but receives a different MAID, it detects a mismerge.

– If the tester receives a CCM with same MD level and correct MAID, but with wrong MEPID, then it detects an unexpected MEPID.

– If the tester receives a CCM with the correct MD, MAID, and MEPID, but with the period value field differs from the tester's CCM interval, it detects an unexpected period.

If the tester misses three CCMs, it declares Loss of Continuity (LoC), which is cleared when the tester detects two consecutive CCMs. LoC can be used to issue traps to the management system, to update the alarm log, and optionally to initiate a switchover to a protection link.

Loopback

Figure 8 shows that Loopback tests are necessary when conducting connectivity and diagnostic tests. The latter includes verifying bandwidth throughput or detecting bit errors. For example, the user can send a Loopback Message (LBM) as a single event or repetitively. The two types of Loopback categories are unicast and multicast.

Loopback tests can be performed both in-service and out-of-service. During in-service tests, the loopback test is performed in the presence of user traffic. The in-service loopback only applies to a configured Ethernet Virtual Circuit (EVC). For instance, the LBM causes only the configured EVC to loop back at the demarcation device; while the other EVC/VLAN continues to pass through the demarcation device.

One application for the LBM is providing connectivity to a remote demarcation device. A tester can be used to send an LBM to the demarcation device. The demarcation device then provides a Loopback Response (LBR) within a specified period of time. If no response is received within that period of time, it declares a LoC.

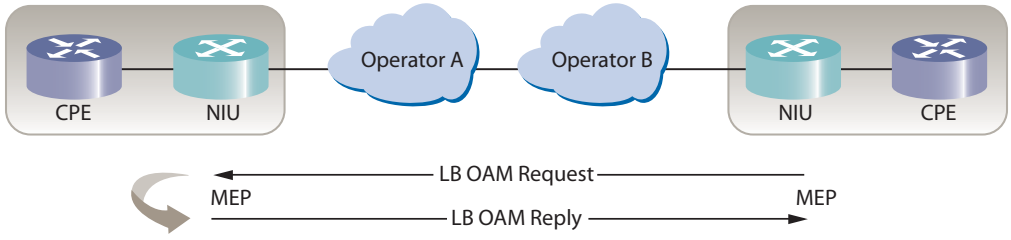

Figure 8: Loopback test

Linktrace

The linktrace traces the path to a target MAC address once a Linktrace Message (LTM) initiates the action. The MIP forwards the message to the destination MAC address where it is no longer forwarded. Each MIP along the path creates a Linktrace Reply (LTR) that is sent back to the originating entity.

Network and Services OAM for Ethernet Networks (Y.1731)

This International Telecommunication Union-Telecom (ITU-T) recommendation was developed in cooperation with the Institute of Electronic and Electrical Engineers (IEEE) 802.1ag, which specifies mechanisms required to operate and maintain the network and service aspects of the Ethernet layer. The recommendation also specifies the Ethernet OAM frame formats as well as the syntax and semantics of OAM frame fields.

Y.1731 defines two sets of functions for fault management and performance monitoring. The fault management functions include many components described in the previous CFM section, such as Continuity Check (ETH-CC), Loopback (ETH-LB), and Linktrace (ETH-LT). The following includes some of the constructs:

– Continuity Check (ETH-CC): As with CCM described in Connectivity Fault Management section.

– Loopback (ETH-LB): Used for both non-intrusive (in-service) and intrusive (out-of-service) applications.

– Link Trace (ETH-LT): As with LTM/LTR described in Connectivity Fault Management section.

– Alarm Indication SignaI (ETH-AIS): Used to suppress alarms upon detection of defect conditions at the server (sub) layer.

– Remote Defect Indication (ETH-RDI): Ethernet Remote Defect Indication function (ETH-RDI): The MEP can use this to communicate to its peer MEPs that a defect condition has been encountered.

– Ethernet Locked Signal (ETH-LCK): Used to communicate the administrative locking of a server (sub) layer MEP and consequential interruption of data traffic forwarded toward the MEP expecting this traffic. It enables the MEP to receive frames containing ETH-LCK information and differentiate it as an administrative locking action at the server (sub) layer MEP rather than a defect condition.

– Ethernet Test Signal (ETH-Test): Used to perform one-way on-demand in-service or out-of-service diagnostics tests, including verification of such things as bandwidth throughput, frame loss, and bit errors. In out-of-service mode, the tester periodically sends ETH-LCK signals up to a user-selectable bandwidth.

– Ethernet Automatic Protection Switching (ETH-APS): Used to control protection switching operations to enhance reliability.

The performance management functions include:

– Frame Loss (ETH-LM): Used to collect counter values applicable for ingress and egress service frames, where the counters tally the number of transmitted and received data frames between a pair of MEPs.

– Frame Delay (ETH-DM): Used for on-demand OAM to measure one- and two-way frame delay as well as frame delay variations.

References

IEEE

802.1D. *Media Access Control Bridges*

802.1Q. *Virtual Bridged Local Area Networks*

802.1ad. *Virtual Bridged Local Area Networks Amendment 4: Provider Bridges*

802.1ah. *Virtual Bridged Local Area Networks Amendment 6: Provider Backbone Bridges*

802.1Qay. *Draft Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks—Amendment: Provider Backbone Bridge Traffic Engineering*

802.1ag. *Virtual Bridged Local Area Networks—Amendment 5: Connectivity Fault Management*

802.3ah. Part 3: *Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks*

P802.1aj/Dx.x. *Virtual Bridged Local Area Networks—Amendment 08: Two-Port Media Access Control (MAC) Relay*

ITU

Y.1731. *OAM functions and mechanisms for Ethernet based network*

G.8010/Y.1306. *Architecture of Ethernet Layer Networks*

G.8110.1/Y.1370.1. *Architecture of Transport MPLS (T-MPLS) Layer Network*

G.8112. *Interfaces for the Transport MPLS (T-MPLS) Hierarchy (TMH)*

G.8121. *Characteristics of Multi-Protocol Label Switched (MPLS) Equipment Functional Blocks*

MEF

MEF 17. *Service OAM Requirements & Framework—Phase 1*

MEF 21. *Abstract Test Suite for UNI Type 2 Part 1 Link OAM*

Glossary

Abbr.	Description
AIS	Alarm Indication Signal
APS	Automatic Protection Switching
CAPEX	Capital Expenses
CB	Customer Bridge
CCM	Continuity Check Message
CE	Customer Equipment
CFM	Connectivity Fault Management
CPE	Customer Premises Equipment
CSU	Channel Service Unit
C-VLAN	Customer VLAN
DA	Destination Address
DEI	Drop Eligibility Indicator
DTE	Data Terminal Equipment
EFM	Ethernet First Mile
E-LMI	Ethernet Local Management Interface
ETH-AIS	Ethernet AIS
ETH-FD	ETHernet Frame Delay
ETH-FL	ETHernet Frame Loss
ETH-LCK	ETHernet LoCK
ETH-LM	ETHernet Loss Measurement
ETH-RDO	Ethernet RDI
EVC	Ethernet Virual Circuit
IEEE	Institute of Electronic and Electrical Engineers
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPTV	Internet Protocol TV
ITU	International Telecommunication Union
LBM	LoopBack Message
LBR	LoopBack Reply
LSP	Label Switched Path
LSR	Label Switch Router
LTM	LinkTrace Message
LTR	LinkTrace Reply
LoC	Loss of Continuity
MA	Maintenance Association
MAC	Media Access Control

Abbr.	Description
MAID	MAID
MC	MultiCast address
MD	Maintenance Domain
MEF	Metro Ethernet Forum
MEN	Metro Ethernet Network
MEP	Maintenance EndPoint
MEPID	MEP ID
MIB	Management Information Base
MIP	Maintenance Intermediate Point
MPLS	Multi Protocol Label Switching
MPLS-TP	MPLS-Transport Profile
MSO	Multiple System Operator
NIU	Network Interface Unit
OAM	Operation, Administration, and Administration
OAMPDU	OAM Protocol Data Unit
OPEX	Operational Expenses
OTN	Optical Transport Networks
PB	Provider Bridge
PBB	Provider Bridged Backbone
PBB-TE	Provider Bridged Backbone - Traffic Enginneering
PBT	Provider Bridged Backbone Transport
PE	Provider Edge
PCP	Priority Code Point
PSN	Packet Switched Network
PM	Performance Monitoring
QoS	Quality of Service
RDI	Remote Defect Indication
RFC	Request For Comments
SA	Source Address
TDM	Time Division Multiplexing
TDMoIP	Time Division Multiplexing over IP
T-MPLS	Transport MPLS
TPMR	Two Port MAC Relay
UNI	User Network Interface
VID	VLAN ID
VLAN	Virtual LAN

Test & Measurement Regional Sales

NORTH AMERICA
TEL: 1 866 228 3762
FAX: +1 301 353 9216

LATIN AMERICA
TEL: +1 954 688 5660
FAX: +1 954 345 4668

ASIA PACIFIC
TEL: +852 2892 0990
FAX: +852 2892 0770

EMEA
TEL: +49 7121 86 2222
FAX: +49 7121 86 1222

www.jdsu.com/test

 30162593 001 0109 ETHOAM.WP.ACC.TM.AE January 2009