

Understanding Ethernet OAM

By Hammadoun Dicko, Product Specialist, EXFO

Service providers and carriers around the world spend billions of dollars to build their networks. As networks become increasingly sophisticated, effective management capabilities become a necessity. Reliability, availability, fast fail over and recovery are all very important parameters that service providers and carriers rely on to ensure service level agreements (SLA).

One of the biggest challenges facing service providers as they deploy Ethernet-based solutions lies in achieving the same level of operations, administration and maintenance (OAM) support required in traditional carrier networks. Essentially, SLAs for specific services need to be met regardless of the underlying technologies used to provide them. Ethernet OAM is one part of the capability needed to meet these SLAs.

WHAT IS ETHERNET OAM?

OAM is a set of functions that provides system or network fault indication, performance monitoring, security management, diagnostic functions, configuration and user provisioning. The purpose of these management tools or capabilities is to enable the monitoring and quick restoration of a network in case of failure. Given that a network is typically comprised of equipment owned by different operators and built by many different manufacturers, OAM has to be standardized to ensure consistency and interoperability. OAM entities are network-aware in that they use information from and provide information to other network entities. They cooperate to provide the consistency and conformity that are critical to an entity's OAM functions.

Ethernet OAM injects OAM packets into the normal stream of data packets at layer 2, and uses end points to process those packets to determine performance, with parameters such as improperly configured nodes, unidentified and out-of-place nodes, disconnected or failed nodes, frame loss, frame delay, end-to-end path identification and bit error rates.

As for Ethernet Service OAM (SOAM), it provides management solutions to guarantee and measure end-to-end performance. It enables end-to-end SLA for standardized Ethernet services, in-service SLA verification as well as network monitoring and troubleshooting from the central office.

Ethernet Service OAM protocols support two sets of functions:

- **Connectivity Fault Management (CFM)**, which detects, verifies and isolates connectivity failure as defined in ITU Y.1731, IEEE 802.1ag and MEF 30.1. This is performed end-to-end, although some functions can isolate faults in segments.
- **Performance Monitoring (PM)**, which provides the capability for performance monitoring as defined in ITU Y.1731 and MEF 35. This is performed end-to-end.

Other OAM protocols include IEEE 802.1ab for link-layer discovery, IEEE 802.3ah for Ethernet in the first mile and ITU-T G.8113.1 for MPLS-TP (multiprotocol label switching – transport profile) OAM. The Metro Ethernet Forum has defined additional service OAM requirements, namely MEF 17.

DEFINITIONS

Maintenance domain (MD): The portion of a network typically owned and operated by a single entity, over which connectivity faults can be managed. MDs are configured with names and eight levels, ranging from 0 to 7. A hierarchical relationship based on these levels exists between domains.

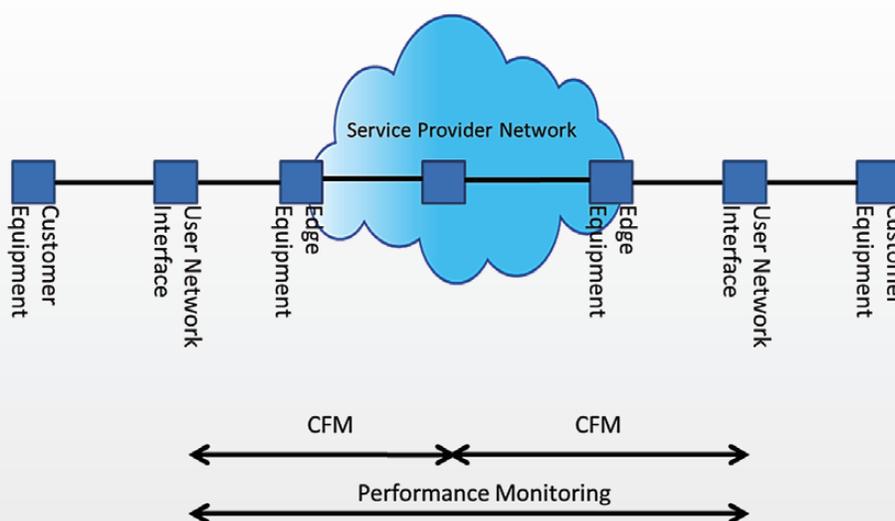


Figure 1. Maintenance domain

Maintenance entity group end point (MEP): The boundary points of a maintenance domain. They can initiate and terminate OAM frames. End-to-end tests are initiated and terminated by MEPs.

Maintenance entity group intermediate point (MIP): The intermediate points in a maintenance domain. They do not initiate OAM frames, but can respond to some OAM frames (loopback and link trace) to isolate faults.

Maintenance entity (ME): This entity requires management and defines a relationship between two maintenance entity group end points.

Maintenance entity group (MEG): A group of MEs that are in the same administrative boundary, with the same MEG level and belonging to the same point-to-point or multipoint Ethernet connection.

Maintenance association (MA): A set of MEPs established to verify the integrity of a single service instance.

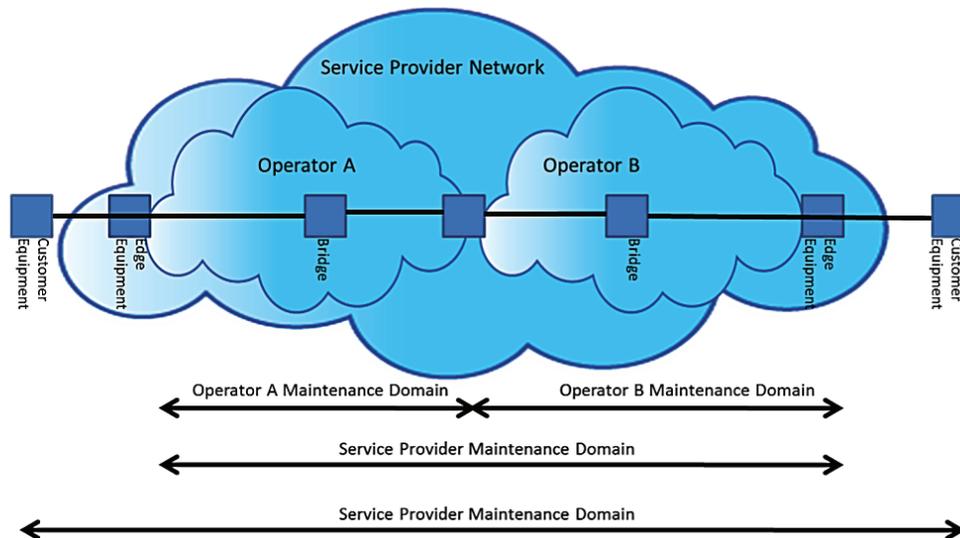


Figure 2. Multiple maintenance domains

SOAM functions are performed end-to-end (i.e., customer to customer). Therefore, when a connectivity fault occurs, it is possible to locate it. In figure 2, a fault can be located in any one of three possible segments. In the service provider's maintenance domain, SOAM functions are performed end-to-end within the MEG, which comprises two MEPs and two MIPs. There too, a fault can be located in any one of the three possible segments. The other two segments belong to operators A and B respectively. Therefore, operators A and B can use S-OAM functions, but only within their respective MEGs.

ETHERNET OAM STANDARDS

This section covers the standards that are available for Ethernet OAM including their different functionalities and use cases.

IEEE 802.1ag

802.1ag focuses on the end-to-end connectivity and continuity of nodes in an Ethernet network. This is why it is referred to as connectivity fault management, or CFM. Since it applies to bridges and bridge applications, it specifies a lot of multicast packets in addition to unicast packets. It also handles both multipoint and point-to-point connections.

802.1ag has three main functions: continuity check messages (CCM), loopback messages and responses (LBM and LBR) and link trace messages and Responses (LTM and LTR).

CCM

A continuity check message is an OAM Protocol Data Unit (OAMPDU) that provides service monitoring from one end point to another (MEP to MEP). The CCMs exchanged between MEPs are analogous to a "heartbeat" message and can be configured to be sent at one of seven standard intervals: 3.3ms, 10ms, 100ms, 1s, 10s, 1 min, and 10min.

CCMs can be either multicast or unicast, although the use of multicast packets is preferred, and they run continuously until turned off. When the reception of CCM messages at a node is lost, this represents the loss of connectivity to that node. The reception of CCM messages from an unknown node represents a possible misconfiguration of nodes.

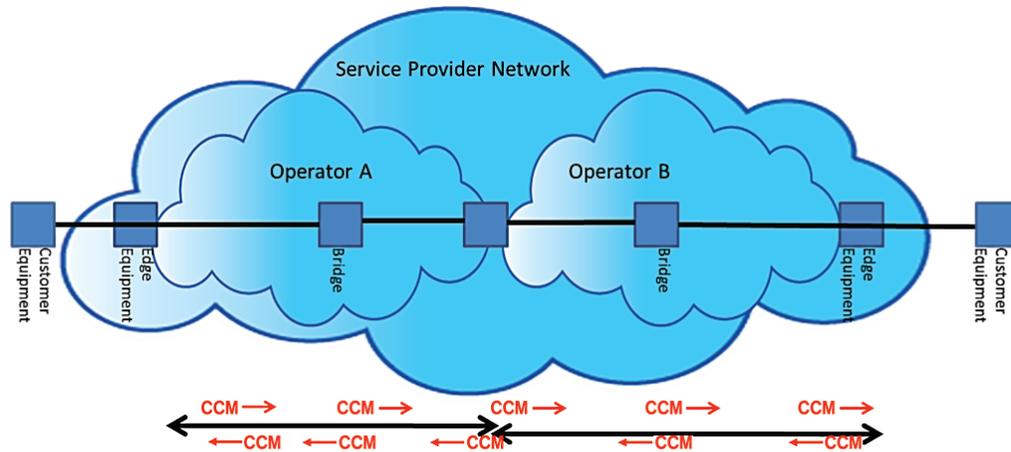


Figure 3. CCM process mechanism

LBM/LBR

LBM and LBR are used to determine the integrity of the data path. Once a network fault has been established, the service provider can verify the loss of service by initiating an 802.1ag loopback test. A loopback is similar to a layer 3 ping request/reply. In a loopback test, a MEP sends messages to another MEP or MIP to verify connectivity across a given MA. These messages can be unicast or multicast, although the use of unicast is preferred.

LTM/LTR

LTM and LTR are used by a MEP to verify the complete link trace to a peer MEP. Each MIP and the peer MEP will respond to an LTM.

Once a network fault has been confirmed by the service provider, the Link Trace feature can be used to isolate its specific location. This feature traces the service from one MEP to another MEP or MIP using its MAC address, and to all the MIPs along the MA.

ITU-T Y.1731

ITU-T Y.1731 is a protocol used mainly for fault management and performance monitoring. It defines performance monitoring measurements (i.e., frame loss ratio, frame delay, frame delay variation, service availability) to assist with SLA assurance and capacity planning. It applies to both multipoint and point-to-point connections and relies on the 802.1ag protocol for transport, making it a type of extension to 802.1ag.

Y.1731 measures the following performance parameters: frame loss, delay, delay variation and service availability. Its main features are alarm indication signals (ETH-AIS), remote defect indication (ETH-RDI), locked signal (ETH-LCK), test signal (ETH-Test), performance monitoring (ETH-PM), frame loss measurement (ETH-LM), frame delay measurement (ETH-DM) and client signal fail (ETH-CSF).

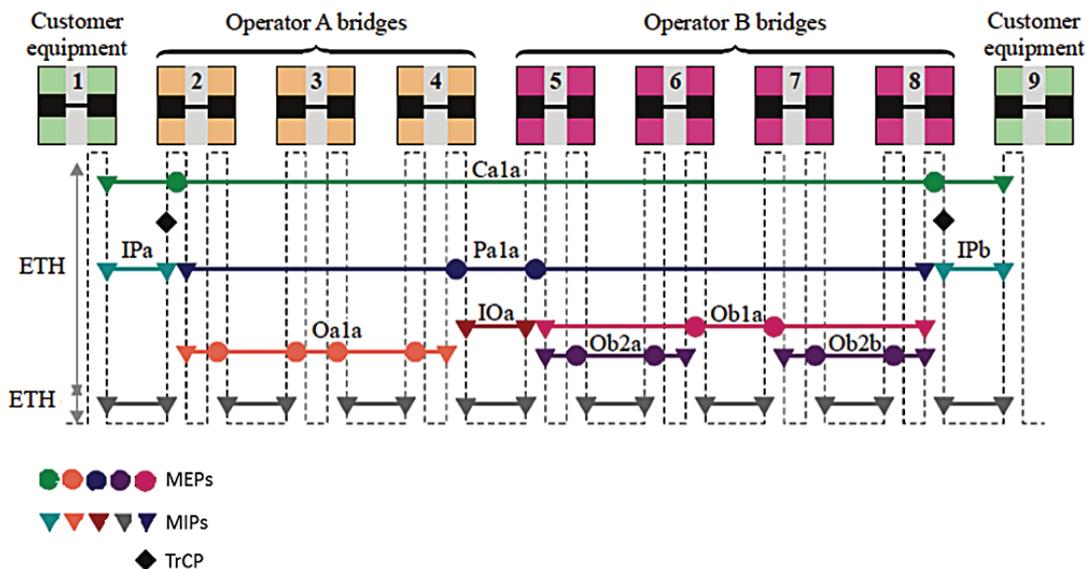


Figure 4. OAM with ITU-T Y.1731

ETH-AIS

This message is sent to the far end when the near end detects an alarm via the CCM packets. The AIS frame is transmitted periodically until the fault condition is resolved.

ETH-RDI

This indication is used for fault notification. If a MEP is defective, it will transmit an RDI to its peer MEPs to inform them that a defect condition has been encountered.

ETH-LCK

This signal is used to communicate an administrative lock to the far end, resulting in an interruption in data traffic. It tells the far end that the near end is present but unavailable for use. LCK frames are also transmitted periodically until the administrator clears the lock.

ETH-Test

This test signal, which is generated by a MEP, is sent to a peer MEP to verify the integrity of the received test signal from the peer MEP. It is also used for bit error rate (BER) and throughput measurements.

ETH-PM

This functionality is used to monitor the performance of traffic from point-to-point or end-to-end on a given domain.

ETH-LM

This feature is used by a MEP to measure the frame loss with the peer MEP in both directions from a single end point.

ETH-DM

This functionality is used by a MEP to measure the roundtrip delay with the peer MEP. This is also used to measure the delay as well as the delay variation.

ETH-CSF

This feature is used for fault notification. It is used by a MEP to propagate to a peer MEP the detection of a failure or event in an Ethernet client signal. It is used when the client does not support fault detection mechanisms.

IEEE 802.3ah

This standard is also referred to as Ethernet for the First Mile (EFM) OAM. While 802.1ag specifies verification of end-to-end connectivity across multiple domains and hops, 802.3ah specifies verification of point-to-point connectivity across only one hop.

It uses the following mechanisms: discovery, loopback, polling of management information base (MIB) variables, remote failing indication and link monitoring.

Discovery

This is the first phase. This is where the devices in the network are identified along with their OAM capabilities. If the discovery fails for any reason, all other OAM activity is aborted until the discovery can be re-established.

Loopback

An OAM entity can put its remote peer into loopback mode using loopback control payload. This helps the administrator ensure link quality during installation or troubleshooting.

Polling of MIB Variables

IEEE 802.3ah provides read-only access to remote MIB, but limited to specific MIB branches and leaves. This feature is based on the fact that the administrator can also retrieve/reset MIB variables at the far end. These variables are one of the primary sources of OAM information available to the system administrator.

Remote Failure Indication

This mechanism enables the OAM entity to convey the degradation of an Ethernet link to its peers via specific tags in OAM payload.

Ethernet Link Trace

This functionality, which is used for fault isolation, allows a MEP to verify the complete link trace to a peer MEP. Each MIP and its peer will respond to the link trace message.

ITU-T G.8113.1/MPLS-TP OAM

The ITU-T G.8113 standard defines the OAM protocol for MPLS-TP in packet transport networks (PTN). It applies to both point-to-point and point-to-multipoint topologies. It relies on 802.1ag and Y.1731 protocols for transport and can be considered as an extension of both of them.

MPLS-TP is a simplified version of MPLS for transport networks, with some of the MPLS functions having been removed, such as penultimate hop popping (PHP), label-switched paths (LSPs) merge, and equal cost multi-path (ECMP). MPLS-TP does not require MPLS control plane capabilities and enables the management plane to set up LSPs manually.

The functions of OAM for MPLS-TP networks are intended to reduce the operational complexity associated with network performance monitoring and management, fault management and protection switching. One of the goals of MPLS-TP OAM is to provide the tools needed to monitor and manage the network with the same attributes offered by legacy transport technologies.

Two important components of the OAM mechanisms are generic associated channel (G-Ach) and generic alert label (GAL). They allow an operator to send any type of control traffic into a pseudowire (PW) or a label switched path (LSP). G-Ach is used in both PWs and MPLS-TP LSPs, whereas GAL is used in MPLS-TP LSPs to flag G-Ach.

THE IMPORTANCE OF TESTING ETHERNET OAM

Ethernet OAM impacts every aspect of Carrier Ethernet service. It is essential to a network because it enables the automated provisioning, monitoring and fault isolation that makes Carrier Ethernet a truly integrated, scalable and interconnected service.

Network equipment supporting OAM features are being massively deployed in networks carrying Ethernet services because they provide essential functions in those networks. It is imperative to make sure that they are properly configured and delivering the features and functionality they are supposed to. Testing OAM services prior to deployment will therefore help service providers and carriers save time and money.

HOW TO TEST ETHERNET OAM

The test methodology or function used will depend on the OAM feature to validate and the network architecture.

Depending on the OAM standard used, the feature set and functionalities will vary from few to many. This section will highlight the main OAM functionalities that must be tested. The different tests can be grouped in two main categories: fault management and performance monitoring.

The service lifecycle has three parts:

- > Provisioning and turn-up (tested and validated using service activation tests, such as ITU-T Y.1564)
- > Performance monitoring (validated with an OAM test)
- > Fault management (tested with an OAM test)

In order to test and validate services and networks carrying those services, it is necessary not only to test the OAM functionalities of the network elements, but also the network and services themselves during provisioning and service turn-up. Therefore, only a test combining ITU-T Y.1564 and Service OAM can attain these objectives.

Fault Management

Depending on the OAM standard being used, several types of tests can be performed to allow the detection, verification, location and notification of various defect conditions.

Continuity Check

This function is used to confirm the presence of a MEP or test equipment on a network and to verify the presence of peer MEPs. During this test, transmitted frames are either received by a peer MEP using a unicast destination address or by all MEPs in the MEG using a multicast destination address.

Loopback Test

During this test, the tester transmits an LBM payload with a specific sequence number. The peer MEP responds to the LBM with an LBR payload. The test validates each received LBR and reports any invalid LBR, invalid payload and LBR timeout.

Link Trace Test

This test verifies that the complete link trace reaches the peer MEP. The tester will send LTM messages and receive LTR messages.

Test Function

During this phase, the tester generates frames with a specific test pattern and sequence number in order to verify the integrity of the signal received by the peer MEP. This test requires two testers: one at each end.

RDI Test

During this phase, the tester generates a remote defect indicator (RDI) to simulate a defect and validate the reaction and behavior of the peer MEP.

Lock Signal Test

This test is used to generate and detect locked signals. When a lock frame is received, the tester sounds an alarm.

CSF Test

During this phase, the tester generates and detects a client signal fail (CSF) to validate the reception and behavior of the peer MEP.

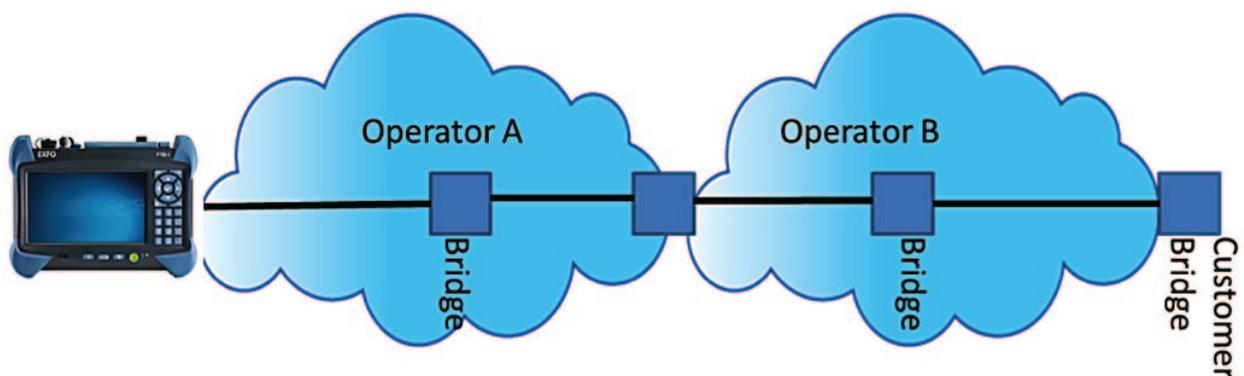


Figure 5: Testing OAM

Performance Monitoring

Performance monitoring is used to measure parameters such as frame delay, frame loss and synthetic loss.

Frame Delay Test

In this test, which is a critical OAM metric, the tester measures the roundtrip delay to the peer MEP. In order to simulate real-world conditions, different frame sizes should be used to validate the frame delay.

Frame Loss

This test checks the bidirectional frame loss occurring with a peer MEP from a single end point. It should be done over as long a period of time as possible to get a good idea of how the network will behave. In order to simulate real-world conditions, different frame sizes should be used to validate the frame loss.

Synthetic Loss

This test uses synthetic frames to check the bidirectional frame loss occurring with a peer MEP from a single end point. This test should also be done using different frame sizes.

CONCLUSION

As multiservice networks get more complex, new technologies will necessarily emerge. Even though the concept of operating, administrating and maintaining networks has existed since the early days of synchronous optical networks and synchronous digital hierarchy (SONET/SDH), Ethernet OAM is continually evolving, with new standards being developed as stakeholders just begin to understand existing ones.

Given the importance that OAM plays in networks, networks must be properly configured to support all the features offered by OAM. However, this can only be achieved by thoroughly testing all the parameters mentioned herein. From a practical standpoint, since the latest technology is not always fully understood by the users, the most efficient approach available to network operators is to use testers equipped with all the required OAM features and metrics.



EXFO Headquarters > Tel.: +1 418 683-0211 | Toll-free: +1 800 663-3936 (USA and Canada) | Fax: +1 418 683-2170 | info@EXFO.com | www.EXFO.com

EXFO serves over 2000 customers in more than 100 countries. To find your local office contact details, please go to www.EXFO.com/contact.