# Ethernet OAM Overview:
# Making Ethernet Manageable

Frank Brockners

Cisco Systems
Hansaallee 249
40549 Düsseldorf
frank.brockners@cisco.com

**Abstract:** This paper provides an overview of three important tools for Ethernet Operations, Administration, and Maintenance (OAM): Link-layer Ethernet OAM, Ethernet-Local Management Interface (E-LMI), and Service Layer Ethernet OAM – commonly referred to as "Connectivity Fault Management".

## Introduction

Over the past years Ethernet evolved from pure local area network (LAN) deployments to metropolitan and wide-area networks. Service providers are leveraging Ethernet to build large and complex networks supporting a wide user base. Enterprises, universities, and governmental institutions connect their different local area networks using Ethernet to establish virtual campuses. It is not uncommon that the delivery of end-to-end services to enterprise customers involves multiple cooperating providers. These trends are flanked by the network assuming a business-critical role. Availability and mean time to repair of an Ethernet network can no longer be compromised, giving rise to a new suite of Ethernet OAM tools which put Ethernet on par with classic transport technologies such as SDH.

Ethernet OAM addresses the following challenges:

- Historically, Ethernet LAN networks have been predominantly managed by network-layer protocols, e.g. Simple Network Management Protocol (SNMP), ICMP Echo (or "IP Ping"), "IP Traceroute". All of which assume a properly operating Ethernet layer. Even if available, "native" Ethernet OAM tools (e.g. Ethernet trace-route) were of proprietary nature and neither covered all the requirements of an Ethernet OAM suite, nor provided interoperability across different vendors.

- An overlay IP infrastructure for management and troubleshooting of Ethernet services may not be feasible for operational or regulatory reasons.

- Leveraging IP OAM tools to manage Ethernet services lacks the per-customer or per-service granularity that is commonly required. IP OAM tools can be employed for fault detection but fail to support fault isolation in the Ethernet network.

- Tools to support automatic provisioning of Ethernet services on the attached customer equipment are typically unavailable – thus requiring fault-prone negotiation and coordination processes between the service provider and the customer.

Ethernet OAM is a broad topic and multiple standard bodies and industry fora tackle the problem. Fortunately the ITU Study Group 13, IEEE 802.3 (clause 57, formerly 802.3ah), IEEE 802.1 (802.1ag Connectivity Fault Management), and the Metro Ethernet Forum (MEF) all drive towards consistent recommendations and standards for Ethernet OAM. We will focus on three main areas of Ethernet OAM which are receiving the most attention in the industry and have shown rapid evolution in the standards bodies: Service Layer OAM (IEEE 802.1ag Connectivity Fault Management), Link Layer OAM (IEEE 802.3ah OAM), and Ethernet Local Management Interface (MEF16 E-LMI). Each of these different OAM protocol suites has unique objectives and complements the others.

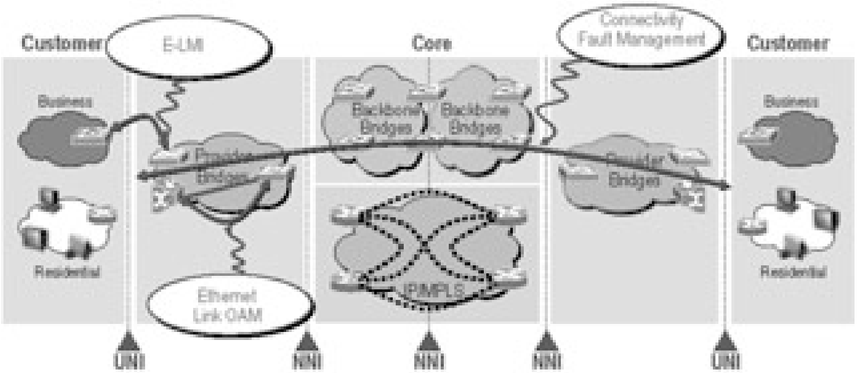

Figure 1: Ethernet OAM - Protocols and where they play

Loosely speaking *Connectivity Fault Management (CFM)* could be understood as “per VLAN” OAM. In other words, it allows service providers to manage each customer service instance individually. A customer service instance or Ethernet Virtual Connection (EVC) is the service that is sold to a customer and is designated by the Service-VLAN tag. The Metro Ethernet Forum defines an EVC as “the architecture construct that supports the association of UNI (User to Network Interface) reference points for the purpose of delivering an Ethernet flow between subscriber sites across the Metro Ethernet Network” [MEF4]. Hence, CFM operates on a per-Service-VLAN (or per-EVC) basis. It enables the service provider to know if an EVC has failed, and if so, provides the tools to rapidly isolate the failure.

Some example scenarios where CFM is relevant:

- A SNMP trap indicates a fault has occurred in the network. How does the service provider identify the set of affected customers, particularly if there are complex failover mechanisms in place?

- How can the service provider proactively discover the failure of a service instance? How can the root cause of the failure (e.g. the failed device) be identified?

- A customer reports partial connectivity within his service instance. How does one identify the failed devices, ports, or even service instances which other operations provide?

- A new customer service instance has just been configured and enabled. How does the service provider confirm that it is operational?

CFM focuses on enhancing the customer experience of a service (i.e. enabling rapid problem identification and isolation so that the appropriate counter actions can be initiated promptly) and is thus a critical tool for operators. Link layer OAM as well as automated customer equipment configuration complement CFM.

*Ethernet link layer OAM* is defined in IEEE 802.3, clause 57. It enables operators to monitor and troubleshoot a single Ethernet link. Therefore it is frequently referred to as “per-hop OAM”. Link-layer OAM was initially developed as part of the “Ethernet in the First Mile” project of the IEEE and thus, it is often referred to by the former project title “802.3ah” but has since enjoyed broad application. Ethernet link layer OAM enables the operator to monitor a link for critical events. In case of issues “loopback” mode can be employed to further testing and problem isolation. Link layer OAM also discovers unidirectional forwarding behavior on links, which occurs when only one direction of transmission fails. Prior to the advent of “IEEE 802.3ah” Ethernet did not provide physical, link-level management.

*Ethernet Local Management Interface (E-LMI)* protocol was developed and ratified by the Metro Ethernet Forum (MEF) as recommendation MEF 16 [MEF16]. E-LMI shares many principles and concepts of the Local Management Interface (LMI) of Frame Relay. E-LMI benefits both the service provider and the end customer. E-LMI operates between the customer edge (CE) device and the user-facing provider edge device (U-PE). Similar to its counterpart in Frame Relay, E-LMI enables the service provider to automatically configure the CE device to match the subscribed service. Thus, the CE device will automatically receive a VLAN-to-EVC mapping and the corresponding quality of service (QoS) settings. Error-prone manual configuration is omitted.

The automatic provisioning of the CE device not only reduces the effort to set up the service, but also reduces the amount of coordination required between the service provider and the enterprise customer. Furthermore, the enterprise customer does not have to learn how to configure the CE device, reducing barriers to adoption and greatly decreasing the risk of human error.

In addition to automatic provisioning of the CE device, E-LMI can provide EVC status information to the CE device. Thus, if an EVC fault is detected (by CFM) the service provider edge device can notify the CE device of the failure which the CE device could use to trigger a corrective action, e.g. switch to a backup link.

### OAM Domain Concept

OAM within service provider networks typically relies on a functional model consisting of hierarchical maintenance domains. Similar to other technologies (e.g. SDH OAM), Ethernet OAM also adopts the maintenance domain model. A maintenance domain is an administrative grouping of devices for the purpose of managing and administering a network. A domain is assigned a unique maintenance level by the administrator, which defines the hierarchical relationship of domains. Maintenance domains may nest or touch, but cannot intersect. If two domains nest, the outer domain must have a higher maintenance level than the one it engulfs. A maintenance domain is constituted by a set of OAM-aware control points (which are typically ports of the involved devices). Typically, only those maintenance points which are located at an edge of a domain (often called maintenance domain end points) are visible to peering operators. Hence, maintenance points within a domain ("maintenance domain intermediate points") are only visible to the operator of the maintenance domain and become invisible at higher maintenance levels. A maintenance domain end point at a lower maintenance level could be a maintenance domain intermediate point at the next level up. The concept of maintenance domains is important due to the different scopes of management that must be provided for different organizations. Often, there are three or more organizations involved in an Ethernet service. Customers purchase Ethernet service from service providers. Service providers may use their own networks, or the networks of other operators to provide connectivity for the requested service. Customers themselves may be service providers, for example, a customer may be an Internet service provider that sells Internet connectivity. Figure 2 illustrates an overview of the OAM domain concept.

Nesting of maintenance domains is especially useful when a service provider establishes agreements with other operators to provide an Ethernet service to a customer. Each operator would have its own maintenance domain, and, in addition, the service provider would define its own domain that would be a superset of the operators' domains. Furthermore, the customer would employ its own end-to-end domain, which, in turn, is a superset of the service provider's domain. Maintenance levels of various nesting domains need to be agreed upon between the involved administering organizations.

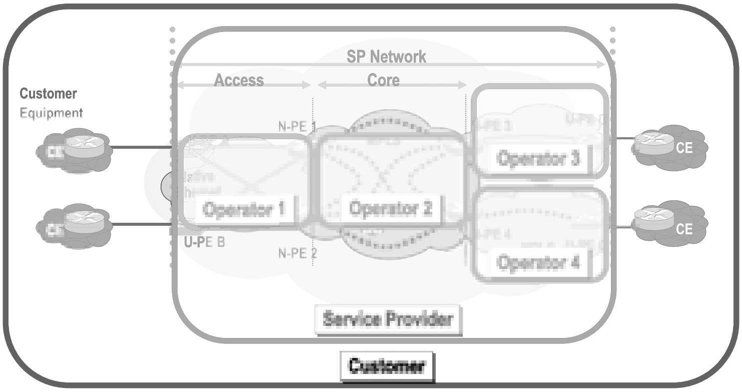

Figure 2: OAM Domain Concept

### OAM Layering Concepts

Although Ethernet OAM provides fault isolation and troubleshooting capabilities for Ethernet services, it does not obviate the need for other OAM mechanisms at other network layers. For example, Ethernet CFM may isolate a fault to a MPLS-based pseudo-wire between two network-facing provider edge (N-PE) devices. However, to determine exactly where the fault has occurred within the MPLS core requires MPLS OAM. MPLS OAM offers similar mechanisms to IEEE 802.1ag: Virtual Circuit Connectivity Verification (VCCV), Bi-directional Forwarding Detection (BFD), LSP ping, and trace-route, which allow the service provider to isolate the fault within the MPLS core. Thus, OAM at each layer in the network helps isolate problems to that layer, and troubleshooting can then be focused on the layer having the problem.

## Link layer OAM: Overview of IEEE 802.3, clause 57

This section discusses the different facets of link-level Ethernet OAM as specified in IEEE 802.3ah-2004, clause 57. Link layer OAM can be implemented on any full-duplex point-to-point or emulated point-to-point Ethernet link. The OAM frames (OAM Protocol Data Units or OAMPDUs) cannot propagate beyond a single hop within an Ethernet network and have modest bandwidth requirements (frame transmission rate is limited to a maximum of 10 frames per second). The major features covered by this protocol are: Discovery, link monitoring, remote fault detection, and remote loopback.

*Discovery*: Discovery is the first phase of link layer OAM. It identifies the devices at each end of the link along with their OAM capabilities.

*Link Monitoring*: Link monitoring OAM serves for detecting and indicating link faults under a variety of conditions. It provides statistics on the number of frame errors (or percent of frames that have errors) as well as the number of coding symbol errors.

*Remote Failure Indication*: Faults in link connectivity that are caused by slowly deteriorating quality are rather difficult to detect. Link OAM provides a mechanism for an OAM entity to convey such failure conditions to its peer via specific flags in the OAMPDUs. Those conditions are a unidirectional loss of signal, an unrecoverable error (such as power failures) or some other critical event.

*Remote Loopback:* An OAM entity can put its remote peer into loopback mode using the loopback control OAMPDU. In loopback mode, every frame received is transmitted back unchanged on the same port (except for OAMPDUs, which are needed to maintain the OAM session). This helps the administrator ensure the quality of links during installation or troubleshooting. This feature can be configured such that the service provider device can put the customer device into loopback mode, but not conversely.

## Overview of Ethernet Local Management Interface (E-LMI)

E-LMI defines the protocol and procedures that convey the information to allow auto-configuration of the customer edge (CE) device by the service provider's user-facing provider edge (U-PE) device. The E-LMI protocol also provides the means for notification of the status of an EVC.

In particular, the E-LMI protocol includes the following procedures:

1. Notification to the CE device of the addition of an EVC. Let us consider the case of a new branch office that connects to the headquarters of a corporation. With the use of E-LMI at the UNIs, the respective CE devices are informed of the availability of a new EVC once the service provider activates the service. In particular, the service end points are notified of the corresponding VLAN ID to be used by a given service (a.k.a. C-VLAN to EVC map attribute)

2. Notification to the CE device of the deletion of an EVC. This is very similar to the previous examples, except the EVC is being removed.

3. Notification to the CE device of the availability (active/partially active) or unavailability (inactive) state of a configured EVC. The primary benefit is that the CE device can take some corrective action, such as rerouting traffic to a different EVC or other WAN service, when informed that an EVC has become inactive.

4. Notification to the CE device of the availability of the Remote UNI. As in the previous case, the CE device can take some corrective action, such as rerouting traffic to a different EVC or other WAN service, when informed that the remote UNI is down.

5. Communication of UNI and EVC attributes to the CE device:

- *EVC identification*: The network informs the CE device as to which VLAN ID is used to identify each EVC (C-VLAN to EVC map). This removes the possibility of a VLAN mismatch between the service provider and customer equipment.

- *Remote UNI identification:* The network informs the CE device of the names of the remote UNIs associated with a given service. This can be used to confirm that the right end points have been connected by an EVC.

- *Bandwidth profiles*: The network informs the CE device of the bandwidth settings of the remote UNI associated with a given service. This allows the CE device to automatically configure its egress traffic shape rate to match the ingress settings of the service provider. Traffic drops at the UNI due to policing of out of profile traffic are avoided and the overall throughput and customer experience are enhanced.

## Service OAM: Ethernet Connectivity Fault Management (CFM)

The term Ethernet “Connectivity Fault Management” (CFM) was coined by the IEEE for their project 802.1ag. CFM constitutes the 5th amendment to Virtual Bridged Local Area Networks (IEEE 802.1Q) and was published in December 2007. The ITU-T counterpart is available as Y.1731 titled “OAM functions and mechanisms for Ethernet based networks” [ITU-Y1731]. Work by IEEE and ITU was conducted in close cooperation between the two organizations.

CFM follows the already described principles of non-intersecting maintenance domains which can be hierarchically nested. Each of these so called maintenance associations is assigned a maintenance level (with a maximum of 8 levels being supported by the standard) to facilitate the hierarchical nesting. CFM aware control points within a maintenance association are referred to as “maintenance association points”. Any port of a bridge could be a maintenance association point – which is typically a configured function of the bridge port. A maintenance point may be classified as a maintenance association end point (MEP), maintenance association intermediate point (MIP), or a transparent point for a maintenance level, in which case it is invisible to CFM operations. Figure 3 shows an example of an Ethernet domain with three different maintenance levels and corresponding maintenance association intermediate and end points. It can be observed that MIPs at a lower maintenance level become MEPs at the next level up. In the example shown, the customer only sees the two UNIs of the service provider. The internal details of the service domain are hidden from the customer.

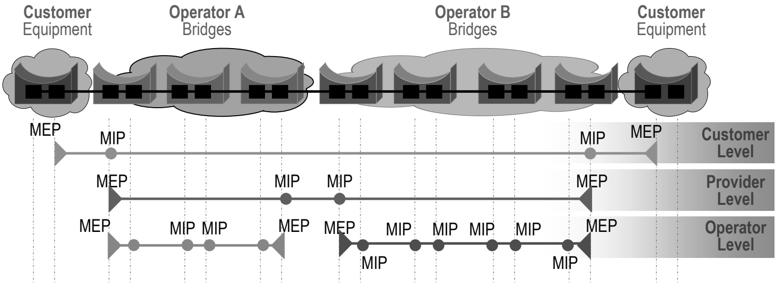

Figure 3: Ethernet OAM domain concept

Maintenance end points reside at the edges of a maintenance domain and are identified by a 13-bit wide MEP-identifier, whereas maintenance intermediate points are internal to the domain. An intermediate point will forward CFM packets (unless it is a loopback or link trace destined for that intermediate point), while end points do not forward CFM packets because they must keep them within the domain. The only exception to this is when an end point is also acting as an intermediate point for a higher-level domain, in which case it will forward CFM packets as long as they are part of the higher-level domain.

Following the requirement for proper layering, CFM uses standard Ethernet frames which are identified as CFM frames by a specific Ether-Type (0x8902). Hence, every bridge, including legacy bridges which do not support CFM, are able to forward CFM messages.

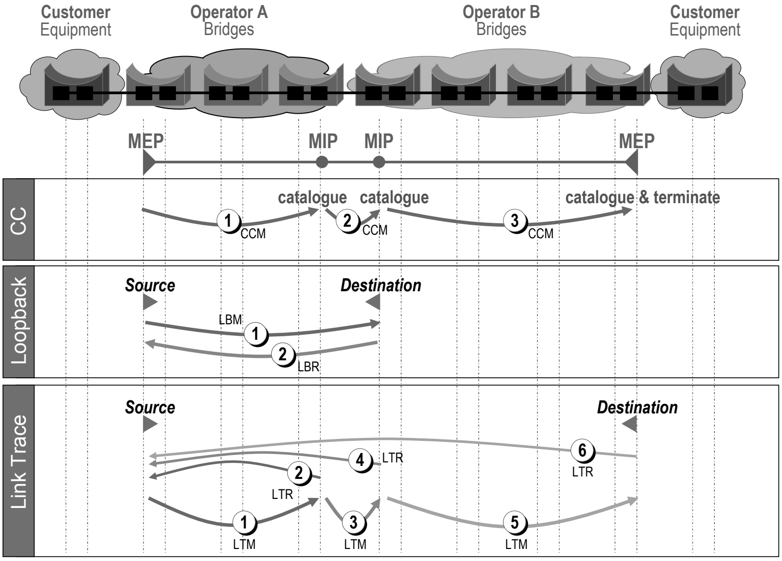

Figure 4: CFM tools: Connectivity Check, Loopback, and Link-Trace

### Connectivity Fault Management Protocols

Ethernet CFM comprises three protocols that work together to help administrators debug Ethernet networks. These are: continuity check, link trace and loopback. In a typical scenario "connectivity check" will make the operator aware of a failure. Loopback can be employed to verify the detected failure and link trace will typically be used to isolate it. If the isolated fault points to a specific link, further OAM tools which are typically specific to the technology of the link can be used for further investigations. In case of MPLS pseudo-wires, tools like Virtual Circuit Connection Verification (VCCV) or LSP ping could be leveraged.

Figure 4 outlines the high-level operation of the three different protocol suites.

- *Continuity Check* messages (IEEE: Connectivity Check Message (CCM); ITU: ETH-CC) are multicast "heartbeat" messages issued periodically by maintenance end points to the group destination MAC address 0x0180C200003y – with y=0,..,7 depending on the level of the domain. Transmission intervals range from 3.3ms to 10min. They allow maintenance end points to detect loss of service connectivity amongst themselves. The default hold time is 2.5-times the transmit interval. They also allow maintenance end points to discover other maintenance end points within a domain, and allow maintenance intermediate points to discover maintenance end points. Note that in networks with a large number of maintenance associations and associated MEPs, CCM can significantly contribute to the load of the network, especially if aggressive intervals are chosen (for details see section 22.5 of [IEEE8021ag]).

- *Link Trace* messages (IEEE: Link Trace Message (LTM) and Link Trace Reply (LTR); ITU: ETH-Trace) are multicast messages sent by a maintenance association end point on the request of the administrator to track the path (hop-by-hop) to a destination maintenance association end point. They allow the transmitting node to discover vital connectivity data about the path. Each traversed MIP/MEP will respond with a unicast link trace reply (LTR) to a LTM.

- *Loopback* messages (IEEE: Loopback Message (LBM) and Loopback Reply (LBR); ITU: ETH-LB) are transmitted by a maintenance end point on the request of the administrator to verify connectivity to a particular maintenance point. Loopback indicates whether the destination is reachable or not; it does not allow hop-by-hop discovery of the path. It is similar in concept to ICMP Echo (Ping). Note that ITU Y.1371 also allows for multicast loopback messages (with the reply always being unicast). Multicast loopback messages can be useful in networks which avoid CC for reasons such as scalability concerns. In those networks, a multicast loopback can play the role of an "on-demand" CC message.

The ITU-T Study Group 13 developed the recommendation Y.1731 in close cooperation with IEEE CFM. Y.1731 acknowledges the CFM principles and further expands the service OAM capabilities. The enhancements are mainly focused on performance monitoring, additional fault notification messages, and the support of alternate protection mechanism (i.e. support of an “Alarm Indication Signal” can provide failure indications in networks which do not employ spanning tree for restoration but rely on 1:1 path protection mechanisms). Performance management parameters are similar to those which have been defined for other packet network technologies such as Frame-Relay and include for example frame loss ratio, frame delay, frame delay variation, throughput or availability. Y.1731 extensions use the same frame format and op-code space as IEEE CFM, facilitating an easy co-existence. It should be noted that the Metro Ethernet Forum also conducts a performance management project which is closely aligned with both the ITU as well as the IEEE efforts and aims at defining performance metrics for point to point and multipoint EVCs (see [MEF17]).

## Acknowledgments

The author would like to gratefully acknowledge people who have contributed to the making of this paper. They include Norman Finn, Aron Lerner, Jose Liste, Eric Matkovich, Sabrina Pittarel, Samer Salam, and Ali Sajassi.

## References

[IEEE8021ag] IEEE P802.1ag, draft 8.1, “Virtual Bridged Local Area Networks – Amendment 5: Connectivity Fault Management.”, June 2007

[MEF 4] Metro Ethernet Forum: MEF 4; “Metro Ethernet Network Architecture Framework - Part 1: Generic Framework”, May 2004

[MEF16] Metro Ethernet Forum: MEF 16; “Ethernet Local Management Interface (ELMI)”, January 2006

[MEF17] Metro Ethernet Forum: MEF 17; “Service OAM Requirements & Framework – Phase 1”, April 2007

[McFarland05] McFarland, M.; Salam, S.; Checker, R.; “Ethernet OAM: Key Enabler for Carrier Class Metro Ethernet Services”, IEEE Communications Magazine, Volume 43, Issue 11, Nov. 2005, Pages: 152 - 157

[Cisco-EOAM] “Overview of Ethernet Operations, Administration, and Management”, Whitepaper, Cisco Systems, September 2006

[ITU-Y1731] ITU-T Recommendation Y.1731 (02/08), “OAM functions and mechanisms for Ethernet based networks”; supersedes Y.1731 (05/06)