Ethernet Operations, Administration, and Maintenance

This paper provides an overview of three important tools for Ethernet Operations, Administration, and Maintenance (OAM) provided by Cisco®, developed in accordance with the leading industry-standards bodies, and how they support the Cisco IP Next-Generation Network (IP NGN) architecture.

Introduction

As intense competition continues to erode their profitability and demand for new services increases, service providers are accelerating their transition to an IP-based NGN. Service providers require innovative, converged infrastructures to improve delivery of current services and provide a scalable framework for tomorrow's new, bandwidth-intensive services such as IPTV, video on demand (VoD), gaming, and voice over IP (VoIP). Solutions that provide greater network intelligence, integration, and flexibility will not only give carriers short-term relief, but also position them to seize new market opportunities. These solutions are part of a larger vision – the Cisco IP NGN – encompassing a broad transformation of not only the service provider's network, but its entire business. The IP NGN empowers service providers to meet the needs of all customer segments efficiently and economically while providing the basis for delivering applications that enable sustainable profitability and subscriber retention.

Convergence is central to the IP NGN, and it occurs in three fundamental ways: application convergence, service convergence, and network convergence. The Cisco IP NGN Carrier Ethernet Design provides the network layer infrastructure of the Cisco IP NGN architecture that is resilient, intelligent, scalable, and oriented toward new service delivery. In order to meet new business challenges, service providers must have sufficient intelligence in the network infrastructure to scale current services and quickly enable new services. The Cisco IP NGN Carrier Ethernet Design is service optimized to deliver both consumer and business services over a single Carrier Ethernet infrastructure. Service Providers will use Ethernet OAM to reach out from the central office (CO) all the way to the customer premise, providing the "eyes, ears, and hands" with which network operations can be performed. Creating a highly reliable and available converged network is a goal that many carriers are already pursuing through their efforts to eliminate multiple service-specific networks or to reduce multiple layers within a network. A "many services, one network" model in which a single network can support all existing and new services will dramatically reduce the total cost of ownership (TCO) for service providers and allow them to quickly provide new services.

Service providers are turning to Ethernet technology for metropolitan-area networks (MANs) and WANs to support the IP NGN model. Ethernet is familiar to enterprise customers and their IT staffs; it can scale to deliver bandwidth up to 10 Gbps to support demanding applications such as triple play, and its bandwidth can be tailored to deliver performance that meets the needs of specific business applications. The challenge for service providers is to provide a highly available Ethernet network where entertainment and business-grade services are assured of being delivered even if the physical link or virtual paths in the network fail. Increasingly, Ethernet as a WAN technology is being recognized as the medium of choice for NGNs. OAM provides the service assurance over a converged network that service providers are looking for in an Ethernet network. Service assurance provides the detection, resiliency, and monitoring capabilities that are needed

for service availability, increased service velocity, allowing auto-provisioning of equipment, and making end-to-end deployment easy through connectivity fault management and link-level protection. Ethernet OAM helps the service provider to provide end-to-end service assurance across the IP/MPLS core, the Ethernet metro, and to the customer's premises.

Overview of OAM

The advent of Ethernet as a metropolitan and wide-area networking technology has accelerated the need for a new set of OAM protocols. Service provider networks are large and complex with a wide user base, and they often involve different operators that must work together to provide end-to-end services to enterprise customers. While enterprise end-customer demands continue to increase, so do the requirements for service provider Ethernet networks, particularly in the areas of availability and mean time to repair (MTTR). Ethernet OAM addresses these challenges and more, thereby directly impacting the competitiveness of the service provider. Ethernet has been used as a LAN technology for many years, and enterprises have managed these networks effectively, primarily with the use of Internet protocols such as Simple Network Management Protocol (SNMP), ICMP Echo (or IP Ping), IP Traceroute, and Cisco Unidirectional Link Detection Protocol (UDLD) and Layer 2 Traceroute (supported in Cisco Catalyst® OS and some Cisco IOS® Software-based platforms). In addition to these troubleshooting protocols, Cisco provides a wealth of other configuration, fault, network management, and performance management tools. Cisco also supports MPLS OAM capabilities such as Virtual Circuit Connectivity Verification (VCCV) and Label Switched Path (LSP) ping on the Carrier Ethernet platforms. To complement these OAM capabilities and to ensure that Ethernet can deliver the required customer service-level agreements (SLAs), Cisco has developed comprehensive Ethernet and IP SLA agents, along with an embedded event manager (EEM), and IPTV video quality tools for automated measurement and troubleshooting of Carrier Ethernet deployments.

Ethernet OAM addresses the following challenges:

- The existing protocols mentioned earlier will not work unless the Ethernet layer is operating properly, making Ethernet OAM a prerequisite.
- Many service providers do not want to overlay an IP infrastructure simply for management and troubleshooting of Layer 2 Ethernet services.
- The current management protocols lack the per-customer or per-service granularity that is required to manage the individual Layer 2 Ethernet services provided to enterprises.
- The existing protocols do not assist with provisioning of Ethernet services, which is particularly difficult when the service provider and end customer must coordinate the configurations on their respective Ethernet equipment.

Ethernet OAM is a broad topic, but this paper will focus on three main areas of Ethernet OAM that are most in need by service providers and are rapidly evolving 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 (MEF-16 E-LMI). Each of these different OAM protocols has unique objectives and is complementary to the others. IEEE 802.1ag Connectivity Fault Management provides "service" management as illustrated in Figure 1.

Figure 1. Elements of Ethernet Service Management

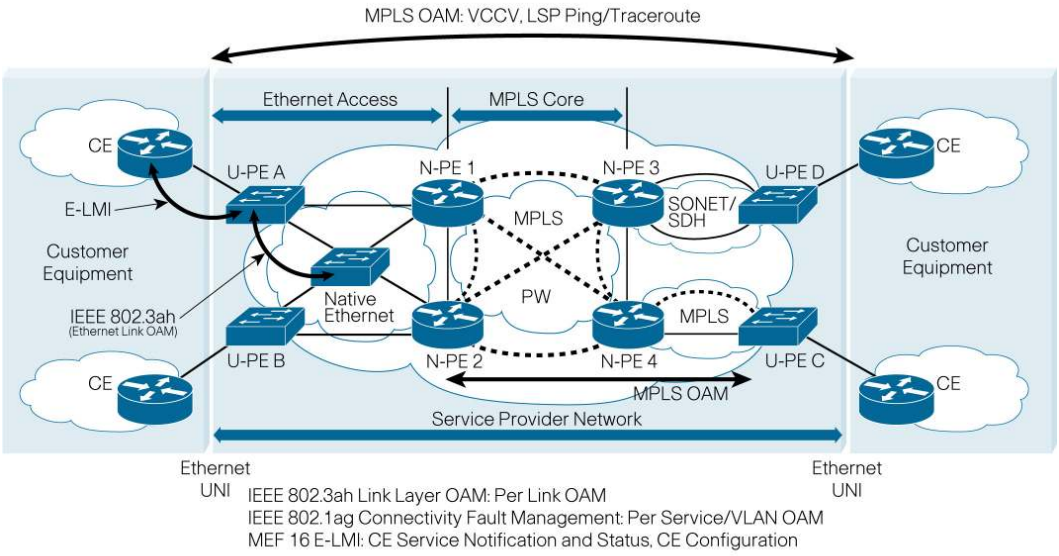

In other words, it allows service providers to manage each customer service instance individually. A customer service instance, or Ethernet Virtual Connection (EVC), is the service that is sold to a customer and is designated by the Service-VLAN tag. Hence, 802.1ag operates on a per-Service-VLAN (or 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.

This functionality is absolutely critical in the following scenarios:

- An SNMP trap indicates a fault has occurred in the network. How does the service provider know exactly which customers are affected, particularly if there are complex failover mechanisms in place?
- An EVC has failed. How does the service provider discover this? And how does the service provider isolate the issue?
- A link or device in an EVC fails. How do the other devices find out so they can reroute around the failure?
- An EVC was just installed. How does the service provider confirm that it is operational?

802.1ag provides the tools to do all of the above easily and quickly, thus reducing operating costs, increasing availability, and decreasing MTTR.

End-to-end service management using 802.1ag is probably the most critical aspect of Ethernet management for a service provider, but another important area is the link management provided by IEEE 802.3ah. Ethernet link management (IEEE 802.3ah) enables service providers to monitor and troubleshoot a single Ethernet link. Even though it was defined for the first-mile connection to the customer demarcation, where most link issues typically occur; IEEE 802.3ah is applicable to any point-point IEEE 802.3 links.

The primary benefits of 802.3ah are that it enables the service provider to monitor a link for critical events and then, if necessary, put the remote device into “loopback” mode in order to do testing on the link. It also discovers unidirectional links, which occur when only one direction of transmission

fails. Current management protocols for Ethernet do not provide the physical, link-level management enabled by 802.3ah.

Ethernet Local Management Interface (E-LMI) protocol was developed and ratified by the Metro Ethernet Forum (MEF) as recommendation MEF 16. E-LMI has substantial benefits to both the service provider as well as the end customer because it brings Ethernet manageability from the service provider network all the way to the customer premises. E-LMI operates between the customer edge (CE) device and the user-facing provider edge (U-PE). Similar to its counterpart in Frame Relay, it 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 bandwidth profile and quality of service (QoS) settings (see Figure 2).

This 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 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 802.1ag), the service provider edge device can notify the CE device of the failure so that traffic can be rerouted to a different path more quickly than if the failure was detected by the routing protocol being run by the CE device. Figure 2 illustrates a high-level overview of Ethernet OAM.

Figure 2. Ethernet OAM Overview

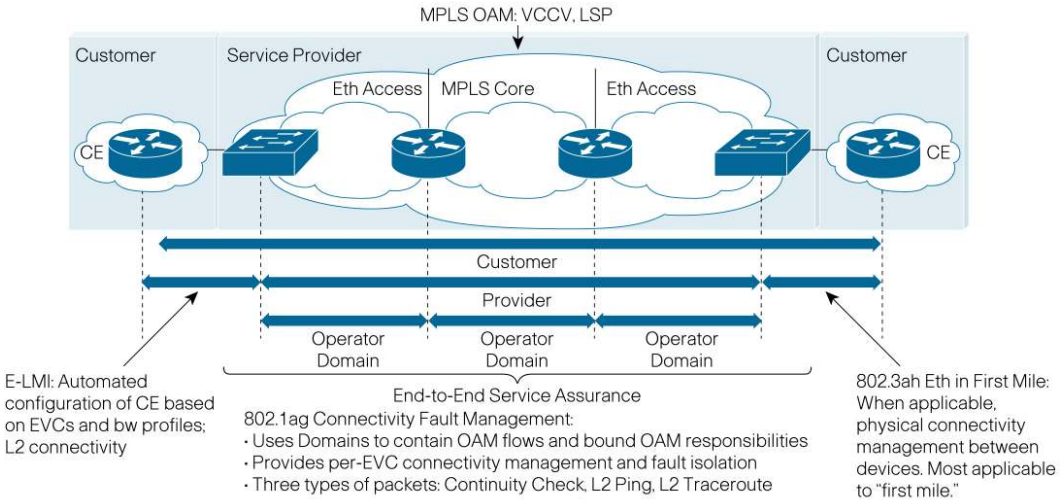

Standards

There is also OAM activity in other standards bodies, particularly the ITU. The primary risk with multiple standards bodies working on the same protocols is that the standards may differ. This can only be mitigated by having the primary authors and editors of the actual drafts present in both standards bodies to achieve consistency. Cisco has led this effort in all of the standards organizations listed in Table 1. For example, the ITU and IEEE have worked on slightly different aspects of a single Ethernet OAM toolset (ITU focusing on requirements and functionality, while IEEE focuses on the protocols themselves).

Table 1. Cisco Involvement in Standards Bodies

Standard	Cisco Involvement	Status
ITU-T SG 13, Y.1731 – Requirements for OAM in Ethernet Networks	Members	Completed
IEEE 802.3ah – Ethernet in First Mile (Physical Link Layer OAM)	Vice Chair and several voting members	Ratified
IEEE 802.1ag – Connectivity Fault Management (Per Service/VLAN OAM)	Co-editor and voting members	In process
MEF 16 – Ethernet Local Management Interface (E-LMI)	Chair of Technical Committee, Editor, voting member	Ratified
IETF – VPLS OAM Requirements and Framework (draft-ietf-l2vpn-oam-req-frmk-01.txt)	Co-editor and voting members	In process

OAM Layers

Although Ethernet OAM provides the fault isolation and troubleshooting capabilities for Ethernet services, it does not obviate the need for other OAM mechanisms at other network layers. For example, 802.1ag Ethernet OAM may isolate a fault to an MPLS-based pseudowire 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 has similar mechanisms to 802.1ag: Virtual Circuit Connectivity Verification, Ping, and Traceroute, 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 problem layer (see Figure 3).

Figure 3. OAM Layers

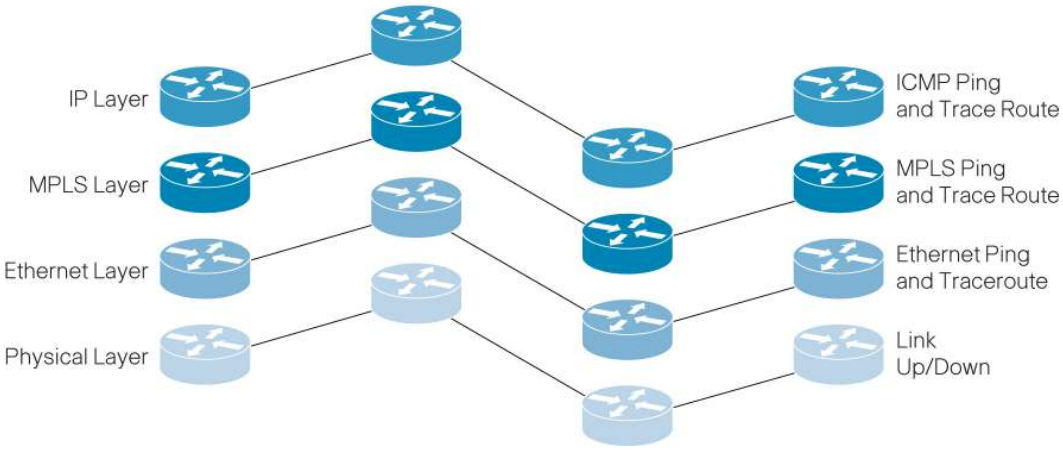

OAM Interworking

Another aspect of OAM is the need for interworking. This is particularly relevant for any type of Alarm Indication Signaling (AIS), in which a node will inform other network nodes that a fault has occurred. Figure 4 illustrates a typical example. In this case, the end customer has some access connections on Ethernet, while others are ATM. MPLS pseudowires are used between the two access networks. When a fault is detected, the AIS must be able to reach all the nodes on the end-to-end connection. Thus, the PE device at the far end of the MPLS core must convert the Ethernet-based AIS into an ATM AIS/Remote Defect Indicator (RDI). The standards for interworking are still in early stages, but this functionality will become an important part of any OAM implementation.

Figure 4. OAM Interworking

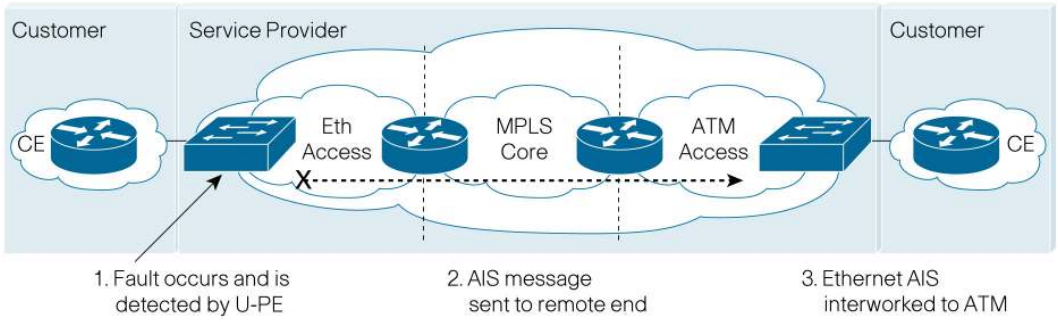

In a typical Carrier Ethernet network, a number of OAM protocols have to work in unison in order to provide end-to-end connectivity monitoring as well as troubleshooting and fault detection capability. OAM Interworking consists of protocol-to-protocol event translations that allow the proper exchange of information among OAM protocols. Consider the following scenarios where the need for Interworking is highlighted:

- Communication of EVC status information (collected via IEEE 802.1ag or MPLS PW OAM) to a CE running E-LMI
- Communication of Link status information (collected via IEEE 802.3ah) to an end-to-end OAM protocol such as IEEE 802.1ag or MPLS PW OAM
- Communication of EVC status information across OAM protocols such as IEEE 802.1ag or MPLS PW OAM

It is imperative that proper protocol layering be honored when defining OAM interworking behavior. Following are some of the fundamentals:

- Strict OAM layering should be maintained: OAM messages should not cross layer boundaries. Each of the service, network and transport layers possesses its well-discernable and native OAM stream.
- OAM messages should not leak outside the confines of a management domain within a layer, where a management domain is governed by a single business organization.
- Interworking is protocol-to-protocol event translation and not necessarily 1:1 message mapping. It is possible that multiple messages at one OAM layer trigger a single message at a client layer.
- Interworking may be inter-layer and/or intra-layer.

Figure 5 shows examples of OAM interworking based on the fundamentals described above.

Figure 5. Examples of OAM Interworking

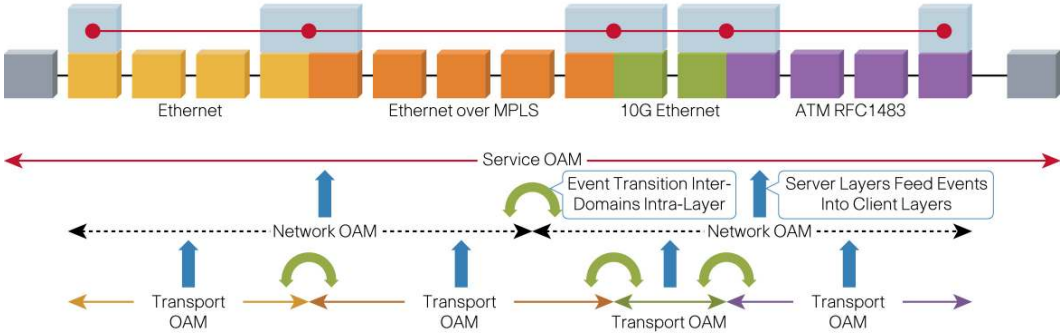

The OAM protocols found in a Carrier Ethernet network will differ by topology, nature of service offering to end customer, device capabilities and provider's preference, but will in general include a subset of: IEEE 802.1ag, E-LMI, IEEE 802.3ah, MPLS Pseudowire OAM and potentially ATM OAM/FR LMI. Given the above list, the following OAM protocol interworking scenarios are possible:

- IEEE 802.1ag → E-LMI
- IEEE 802.3ah → IEEE 802.1ag
- MPLS PW OAM ←→ IEEE 802.1ag

Note the directionality of the arrows (→ or ←→) in the above scenarios. Some of the interworking functions are bidirectional while others are unidirectional. This is mostly dictated by how the constituent protocols are defined to operate.

OAM Integration with NMS/EMS Management Systems

Management systems play an important role in configuring OAM functionality consistently across all devices in the network, and for automating the monitoring and troubleshooting of network faults.

Unlike MPLS OAM where mechanisms such as LSP ping and trace work without any configuration, the mechanisms provided by 802.1ag require configuration before being functional. If you issue an 802.1ag ping or link trace in a network that has not been configured, you will not receive any response. There are two phases to this configuration. The first phase is the network provisioning phase, which enables Connectivity Fault Management (CFM) on the devices and sets up Maintenance Domains (MD) and Maintenance Intermediate Points (MIP). Maintenance Endpoints and Intermediate Points are illustrated in Figure 8. This process needs to be integrated with the network engineering processes of the network operator. Operators are extremely careful about change management for the basic configuration of the network, and it is unlikely they will accept additional tools that manipulate these device configurations, but will in most cases extend their existing tools to include this in the basic configuration that is applied when a device is commissioned. Note, however, that this configuration is a once-only process, and the configuration should not need to be modified again.

The second configuration phase is the service activation phase. Every time a new endpoint is provisioned on a VLAN (service VLAN or customer VLAN), the endpoint needs to be configured as a Maintenance End Point (MEP) – this is to enable the origin of ping and trace packets, as well as configure continuity-check and cross-check functionalities. Because this is a per-service-activation

activity, this needs to be integrated into the normal service-activation flow, and into any provisioning system used for this.

In addition to simplifying configuration, the network management system/element management system (NMS/EMS) simplifies OAM monitoring and troubleshooting. When a problem occurs in the network, the management system will receive messages from the devices and raise alarms accordingly. Most importantly, however, is to have troubleshooting support in the management system. Although the protocols described in this paper enable trouble detection and troubleshooting, in a real situation the troubleshooting steps and decision trees involved require expert knowledge, and manual troubleshooting is prone to human error. Automatic troubleshooting can mimic the actions of an expert and carry out troubleshooting steps faster, hence minimizing service downtime. There is also a security benefit, because less people need to be able to log on to the network devices. A final benefit of automated troubleshooting from an EMS/NMS is in complex cross-domain scenarios, where the Ethernet service is composed of different underlying technologies (for example, an Ethernet connection failure due to an LSP black hole in the MPLS network). It is unlikely that the Ethernet expert is also an MPLS expert, so in a manual troubleshooting scenario, there would be a handoff between experts, or even between operational organizations.

Use Cases

Use Case 1

A service provider may today use SNMP to periodically poll devices for statistics and receive SNMP traps when faults occur. However, when a fault occurs, the service provider has no idea which customers are impacted. By implementing 802.1ag CFM, the continuity-check mechanism will determine which EVCs are impacted so the service provider knows exactly which customer services are down. The operator can verify the loss of connectivity using CFM loopback (ping), and localize the connection failure using CFM link trace. The problem can then be further diagnosed and remedied. Finally, CFM loopback may be used to verify that the remedial action has succeeded and that the service has been re-established.

The use of a management system can automate this process, allowing the operator to automatically:

1. Detect the problem by receiving the CFM traps and correlating the traps from each endpoint into one alarm for the connectivity problem
2. Localize and diagnose the problem
3. Repair any configuration issues, or reroute around the problem
4. Verify that service has been re-established

Obviously the primary advantage of the automation is the faster time to repair and the reduced need for human intervention. It is important to note, however, that 802.1ag is not a resiliency mechanism like Rapid Spanning Tree Protocol (IEEE 802.1w), which provides sub-second failover. Rather, it is a protocol that provides consistent monitoring of EVCs and the troubleshooting tools to isolate faults and verify connectivity.

Use Case 2

A service provider today is likely deploying Carrier Ethernet services without support of a protocol that could automatically convey EVC attribute information to the CE. The most basic problem like a VLAN ID mismatch between customer and service provider may be very involved and require troubleshooting from both entities. With support of E-LMI across devices in a given UNI, EVC information can be propagate to the customer such that the kind of problems mentioned earlier could be completely detected and corrected by the customer without operations involvement providing an immediate OPEX benefit to the service provider

Use Case 3

One of today's challenges faced in the deployment of Carrier Ethernet services is the potential issue that occurs when a failure of an EVC or a remote UNI does not translate into a link status event at the CE. One of the options to handle this scenario is to communicate to the CE the status of an EVC and/or remote UNI with E-LMI between the customer and the service provider network. E-LMI acts as a messenger but does not control or determine the state of a connection. However, E-LMI can learn this information from the knowledge collected by other OAM protocols like IEEE 802.1ag. As such, IEEE 802.1ag CFM to E-LMI OAM Inter-working provides a viable alternative to achieve end-to-end CE fault notification. Upon notification by the network that a given EVC is inactive, the CE can re-route traffic onto another EVC that is acting as the backup path.

Service OAM: Introduction to Connectivity Fault Management (IEEE 802.1AG)

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 protocols.

- **Continuity Check** – These are "heartbeat" messages issued periodically by maintenance endpoints. They allow maintenance endpoints to detect loss of service connectivity amongst themselves. They also allow maintenance endpoints to discover other maintenance endpoints within a domain, and allow maintenance intermediate points to discover maintenance endpoints.
- **Link Trace** – These are transmitted by a maintenance endpoint on the request of the administrator to track the path (hop-by-hop) to a destination maintenance endpoint. They allow the transmitting node to discover vital connectivity data about the path. Link trace is similar in concept to UDP Traceroute.
- **Loopback** – These are transmitted by a maintenance endpoint 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).

Maintenance Domains

Ethernet CFM, within any given service provider network, relies on a functional model consisting of hierarchical *maintenance domains*. A *maintenance domain* is an administrative domain for the purpose of managing and administering a network. A domain is assigned a unique maintenance level (among eight possible) by the administrator, which is useful for defining 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 defined by provisioning which bridge ports are interior to the domain. Figure 6 illustrates this concept.

Figure 6. Ethernet CFM Maintenance Domain

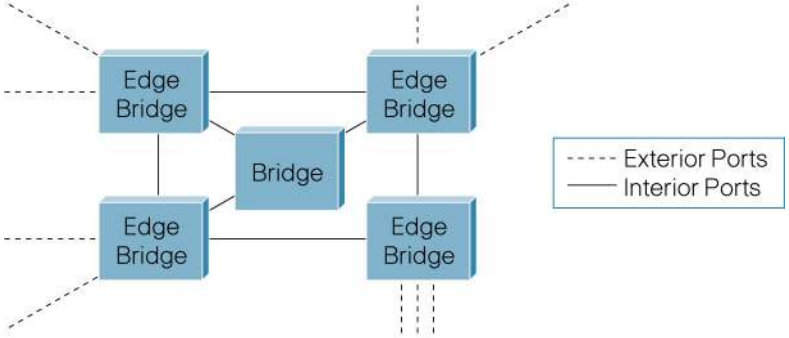

The concept of maintenance domains is important due to the different scopes of management that must be provided for different organizations. Often there are three different organizations involved in a Metro Ethernet service: *customers*, *service providers,* and *operators*. *Customers* purchase Ethernet 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.

Nesting of maintenance domains is useful when considering the business model where the service provider contracts with one or more operators to provide the Ethernet service to a customer. Each operator would have its own maintenance domain, and, in addition, the service provider defines its own domain that is a superset of the operators' domains. Furthermore, the customer has its own end-to-end domain, which, in turn, is a superset of the service provider's domain. Maintenance levels of various nesting domains should be communicated between the involved administering organizations. One model would be where the service provider assigns maintenance levels to the operators. Figure 7 illustrates the logical nesting of domains among organizations.

Figure 7. Hierarchical Nesting of Maintenance Domains

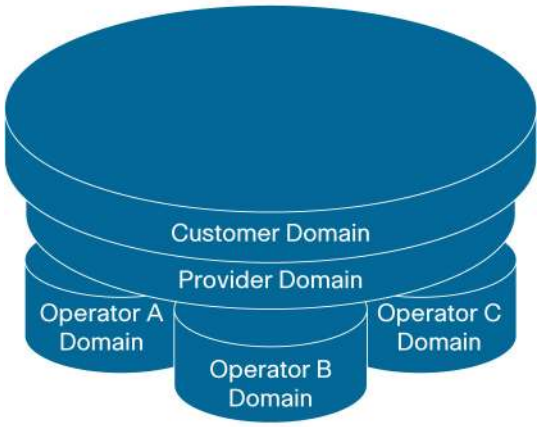

Maintenance Points

Any port of a bridge is referred to as a maintenance point. A maintenance point may be classified as a maintenance endpoint, maintenance intermediate point, or transparent point for a maintenance level (Table 2).

Table 2. Maintenance Point Classifications

Functions	Maintenance Endpoint	Maintenance Intermediate Point	Transparent Point
Initiate CFM messages	Yes	No	No
Respond to loopback and link trace messages	Yes	Yes	No
Catalogue continuity-check information received	Yes	Yes	No
Forward CFM messages	No	Yes	Yes

Maintenance endpoints reside at the edge of a maintenance domain, whereas maintenance intermediate points are internal to the domain. Hence, an intermediate point will forward CFM packets (unless it is a loopback or link trace destined for that intermediate point), while endpoints do not forward CFM packets because they must keep them within the domain. The only exception to this is when an endpoint 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.

Figure 8 shows an example where a service provider is using the networks of two operators to provide service. The **service provider maintenance level** is shown in blue. The maintenance levels for **Operator A** and **Operator B** are shown in orange and violet, respectively. Two special-case maintenance levels are the **customer level** (shown in green) and the **physical layer level** (shown in black). The **customer level** allows the customer to test connectivity (using connectivity checks) and isolate issues (using loopback and link trace). The **physical layer level**, on the other hand, defines the narrowest possible maintenance domain: a single link domain.

Note that the designation of maintenance points as maintenance endpoints or maintenance intermediate points for **Operator A** (or **Operator B**) level is relative to that level only. When these maintenance points are observed relative to the **service provider level**, maintenance endpoints at the **Operator A** level translate into either maintenance endpoints or maintenance intermediate points at the **service provider level**. Furthermore, maintenance intermediate points at the **Operator A** level translate into transparent points at the **service provider level**. Also note that the demarcation of maintenance points as maintenance endpoints or maintenance intermediate points within a domain is left to the discretion of the administrator, because these points mark yardsticks of particular relevance for the management of the network.

Here is an example of how the CFM messages are used across the domains. In Figure 8, the **customer** could use CFM loopback or link trace to isolate a fault between the maintenance point **M** (which is on the CPE) and the intermediate point **L** (which is on the user-facing provider edge equipment, or U-PE). By definition, the link between the CPE and U-PE is a single hop and, therefore, the **customer** would know which link has the fault. However, if the fault is between the two intermediate points (the **L's**), the **customer** will need to rely on the **service provider** to determine between which maintenance (**M**) or intermediate (**L**) points the fault has occurred. Even then, the **service provider** may simply isolate the fault to a specific **operator's** network, and will in turn rely on the **operator** to isolate the fault to a specific link in its network.

Thus, each different organization (**customer**, **service provider**, and **operator**) has the ability to isolate the fault within the organization's maintenance level, without the **service provider** having to

share its network information to the **customer**, or the **operator** having to share its network information to the **service provider**.

Figure 8. Maintenance Points and Maintenance Domains

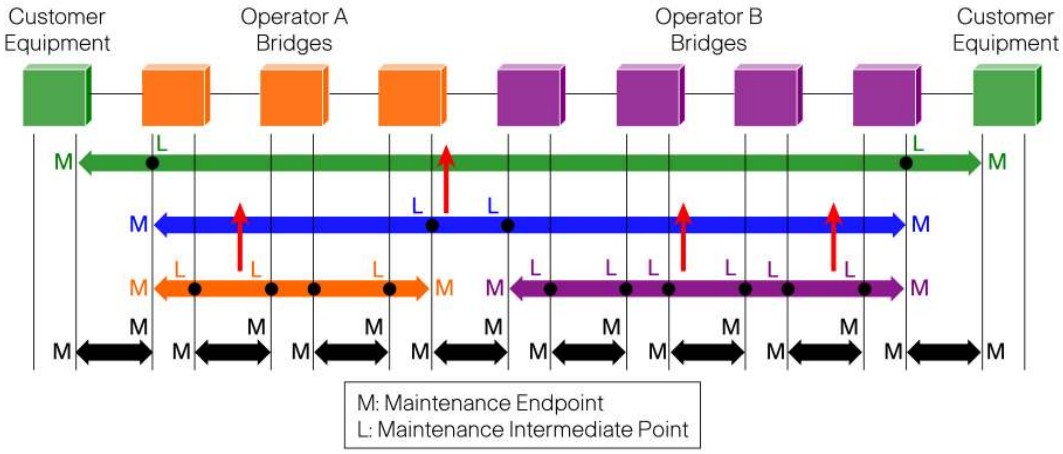

Introduction to Ethernet Local Management Interface

E-LMI defines the protocol and procedures that convey the information that allows auto-configuration of the 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. (See Figure 9.)

Figure 9. Ethernet Local Management Interface

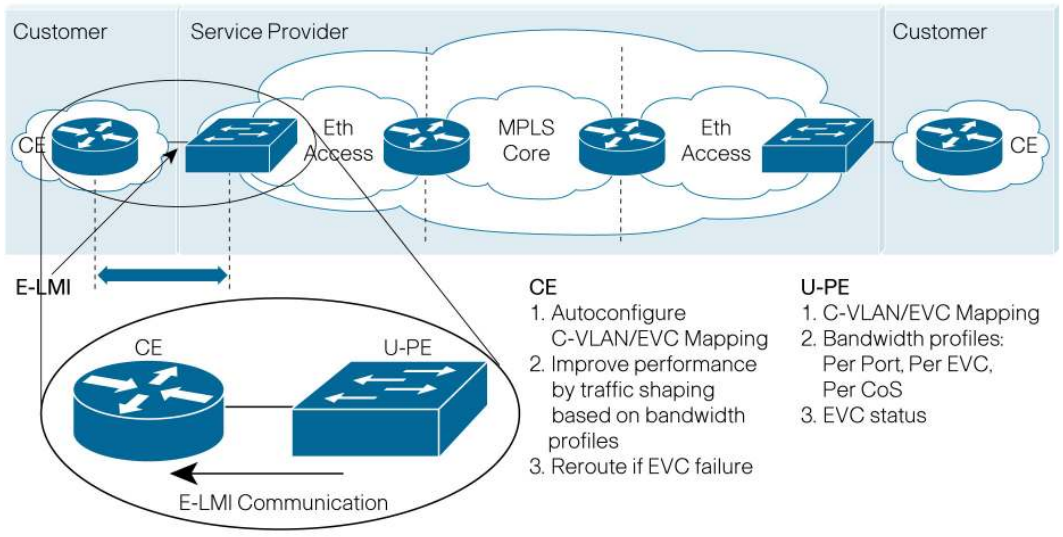

In particular, the E-LMI protocol includes the following procedures:

1. Notification to the CE device of the addition of an EVC:

   An example use case of this is if a new branch office is connected to headquarters. With the use of E-LMI at the UNIs, the respective CPEs are informed of the availability of a new EVC once the Service Provider turns on the service. In particular, the service endpoints 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 SP and customer's equipment.

   - Remote UNI identification

     The network informs the CE device of the names of the remote UNIs associated to a given service. This can be used to confirm that the right endpoints have been connected by an EVC

   - Bandwidth profiles

     The advantage of this is that if the enterprise has subscribed to a 50-Mbps service, then the CE device can automatically configure itself to shape the egress traffic to 50 Mbps on the WAN interface. By shaping to 50 Mbps rather than having the service provider police a 100-Mbps stream down to 50 Mbps, enterprise customers will reduce the number of dropped packets and increase the throughput they receive.

Link OAM: Introduction to 802.3AH OAM

This section discusses the different facets of link-level Ethernet OAM (Figure 10), as specified in IEEE 802.3ah-2004 Clause 57. 802.3ah OAM can be implemented on any full-duplex point-to-point or emulated point-to-point Ethernet link. The 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.

Figure 10. Link OAM – IEEE 802.3ah OAM

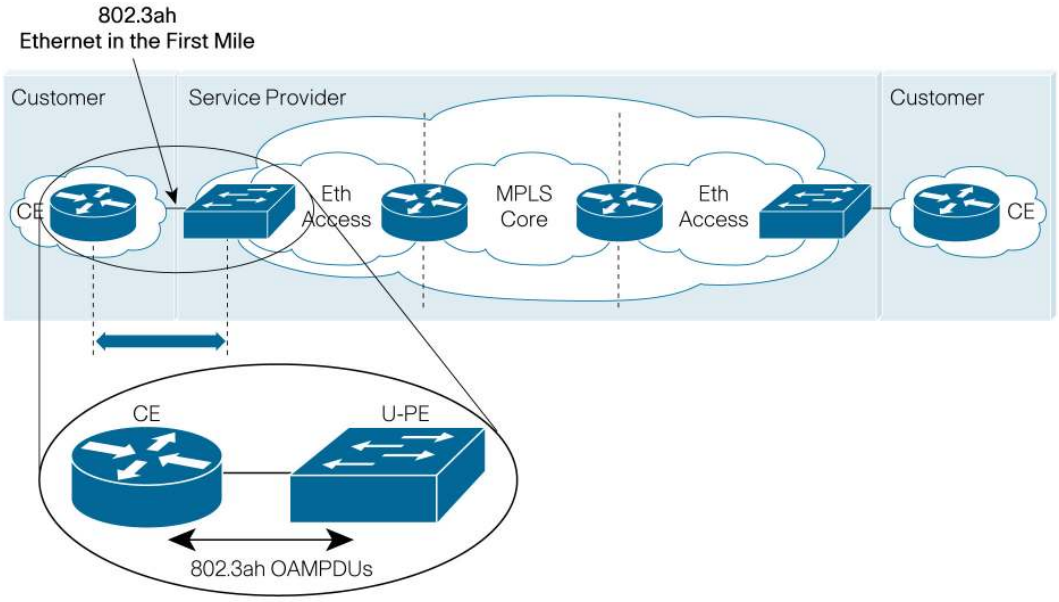

Discovery

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. The failure conditions that can be communicated are a loss of signal in one direction on the link, an unrecoverable error (such as a power failure), or some 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 when troubleshooting. This feature can be configured such that the service provider device can put the customer device into loopback mode, but not conversely.

Summary

At the foundation of an IP NGN are the end-to-end service-assurance capabilities provided by the network layer. Comprising customer element, access/aggregation, intelligent IP/MPLS edge, and multi-service core components with transport and interconnect elements layered below and above, the network layer is also undergoing dramatic and fundamental change compared to only a few years ago. IP/MPLS is being integrated throughout each section of the network along with complementary end-to-end OAM capabilities at the MPLS and Ethernet service layers. Edge and core areas are converging, with each adopting capabilities of the other and providing greater efficiencies to the service provider. Customer elements are converging around Ethernet technology as well. Service providers can take advantage of this convergence to offer newer and more cost-effective Ethernet services while taking advantage of the end-to-end service-assurance capabilities as offered by standardized Ethernet OAM on Cisco routing and switching platforms. Using these capabilities, service providers can deliver both entertainment- and business-grade services over a highly available Ethernet network, to meet and exceed the expectations of their end customers.

Cisco takes a complete lifecycle approach with service providers. This begins with helping the service provider strategize an offering and assisting with the design, plan, development, test, and trial of a service. Cisco also has the marketing resources and expertise to help the service provider position and sell a service, and help match services to customer needs. This level of commitment and support not only distinguishes Cisco from the competition, but also gives Cisco insight into ways to continue to advance technology – insight that goes into its products and solutions. Cisco has introduced innovations in routing, switching, optical transport, security, VoIP, and other technology areas through its participation and leadership in various industries, national, and international standards organizations. Cisco innovations and commitment to open standards benefit the entire networking industry and make it easier for complex, multi-vendor networks to interoperate. The value Cisco delivers through its extensive portfolio of systems offers service providers flexibility in how they choose to deploy these platforms. Cisco combines system architecture, silicon processing, and software services into intelligent networking, allowing these components to be used in multiple different platforms, so service providers can mix and match different systems with confidence that they will all work together.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-350532-02 09/07

