END-TO-END SERVICE OAM

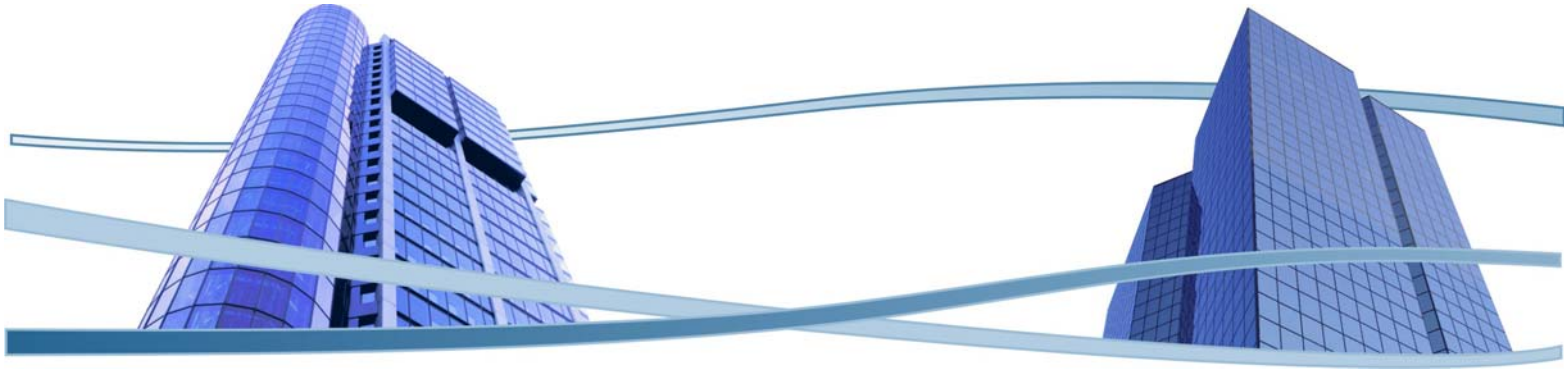

Carrier Ethernet White Paper

An overview of the relevant standards and key capabilities of Ethernet Service OAM in multi-operator networks, and the pivotal role of the Network Interface Device in enabling the delivery of reliable and profitable Ethernet services.

Omnitron Systems Technology, Inc.

140 Technology Dr. Irvine, CA 92618 USA 800-675-8410 +1 949-250-6510 www.omnitron-systems.com

Introduction

Carrier Ethernet has grown to become a dominant technology in Service Provider networks. Driving this growth is the demand from business customers for scalable services, higher bandwidth and lower costs. Carrier Ethernet is defined as the extensions, or enhancements to Ethernet which enable Service Providers to provide dependable Ethernet business services. While traditional, LAN-based Ethernet is a best-effort technology, Carrier Ethernet has evolved to provide robust reliability and availability required for business applications.

As Carrier Ethernet service deployments increase year after year, the number of Ethernet wholesaling partnerships between network Operators is also growing. These partnerships allow Service Providers to link their networks and extend services into new territories that otherwise could not be reached. The underlying transport, metro and access technologies to deliver these services, and the actual mechanisms to interconnect these networks vary from Operator to Operator. This paper assumes that the service has been properly planned and provisioned. Once a service is operational, tools are needed to monitor and ensure that the service meets its contracted Service Level Agreement (SLA). This becomes more important as networks grow to be regional, national, or even global. Also, when the services deteriorate or fail, tools are needed to rapidly isolate, diagnose and repair the failing network segment.

To provide this kind of end-to-end network fault and performance monitoring, Service Providers require a comprehensive set of Operations Administration and Maintenance (OAM) tools. IEEE 802.1ag and ITU-T Y.1731 standards define these Service OAM tools that follow the service path and monitor the entire Ethernet *service* from end-to-end.

With Service OAM, Service Providers can offer Service Level Agreement assurances and reduce operating costs associated with manual network fault monitoring, truck rolls and labor-intensive performance measurements.

The goal of this paper is to summarize the multi-operator and multi-domain network models, and provide an overview of the relevant standards and key capabilities of Ethernet Service OAM. This paper will also discuss the Network Interface Device (NID) and its pivotal role in enabling Carrier Ethernet Service Providers achieve success.

The Carrier Ethernet Ecosystem

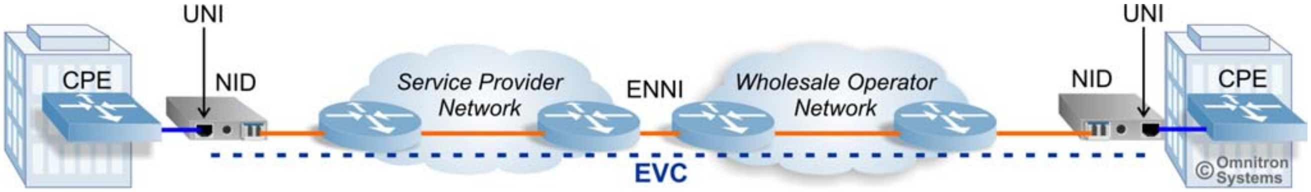

Figure 1. Carrier Ethernet Network Diagram

Figure 1 illustrates a typical point-to-point Ethernet Virtual Connection (EVC) between two Subscriber locations. This Carrier Ethernet network diagram contains the following elements:

CPE: The Customer Premise Equipment is Subscriber-owned network equipment at the edge of the Subscriber's LAN or Enterprise network, and interfaces with a Carrier's Metro Ethernet network.

UNI: User-to-Network Interface is the physical interface or port through which the Service Provider delivers the service to the Subscriber. The UNI is where the point of responsibility changes from the Service Provider to the Subscriber. The UNI can also be considered as the starting and ending point for the EVC. The UNI is typically the customer-facing port on the NID.

NID: The Network Interface Device is demarcation equipment that delivers Ethernet services via a UNI interface. The NID ensures service quality, monitors performance, facilitates troubleshooting, and converts between different physical media (for example, converting the Service Provider's fiber interface to a copper UTP Subscriber interface). The NID is typically located at the Subscriber's site, and is owned and managed by the Service Provider or Operator.

ENNI: The External Network to Network Interface provides the interconnection point between the two networks when more than one Service Provider is involved in delivering an Ethernet service. In the example above, it is the point where the Service Provider and the Wholesale Operator hand off the service(s) to each other.

EVC: The Ethernet Virtual Connection represents the Ethernet service or connection between two or more UNI interfaces. EVCs can be port based (i.e., all Subscriber traffic on a UNI port), or virtual (i.e., all Subscriber traffic carried over a Service Provider VLAN). EVCs can be point-to-point and connect two UNIs, or multipoint-to-multipoint and connect more than two UNIs.

Types of Ethernet Services

The EVC in Figure 1 is a point-to-point E-Line service, and can be used for either Ethernet Private Lines or Ethernet Virtual Private Lines. Ethernet Private Lines use dedicated UNIs for point-to-point connections, and contain a single EVC per UNI. Ethernet Private Lines are the most popular Ethernet service due to their simplicity. Ethernet Virtual Private Lines (EVPL) are often used to replace Frame Relay or ATM services, by supporting a service-multiplexed UNI (multiple EVCs per UNI), and supporting multiple EVPLs from a single physical connection (UNI) at the customer premise.

Other Ethernet service types include E-LAN and E-Tree. E-LAN Service is comprised of multi-point to multi-point EVCs. E-LAN can be used to create a Transparent LAN Service, such as a multi-point L2 VPN. E-Tree is a rooted point to multi-point service and is used for broadcast applications where services flow from a centralized point.

Throughout this paper, the point-to-point EVC (E-Line) is used in the example figures for simplicity.

In Franchise and Out of Franchise

When a Service Provider delivers a Ethernet service to a Subscriber within the Service Provider's own network footprint, this is referred to as In Franchise (also On-Net).

When a Service Provider partners with a another Service Provider such as a regional Operator to deliver Ethernet services to Subscriber locations outside the Service Provider's own network footprint, this is referred to as Out of Franchise (or Off-Net). This requires wholesaling partnerships between Service Providers, and these partnerships allow providers to link their networks and extend services into new territories that otherwise could not be reached.

OAM can be challenging to implement in a single-operator network, and the challenges to deploy OAM tools grow more complex when Out-of-Franchise services span multiple Operator networks.

The Evolution of OAM

OAM tools and functionality have evolved as Ethernet has evolved from the LAN to the robust telecom service it is today.

Ethernet OAM began as simple Remote Monitoring (RMON) that monitors ports and utilization on Ethernet LANs and is limited to port-level frame data.

Standards bodies have defined modern Ethernet OAM tools to provide comprehensive fault management and performance monitoring capabilities that could ensure the same reliability as the famous Five Nines (99.999%) in TDM networks.

Link OAM

The first OAM standard developed for Carrier Ethernet was IEEE 802.3ah, also known as Ethernet in the First Mile (EFM) or Link OAM because it provides fault detection and monitoring of the physical links in first mile access networks. 802.3ah is a good tool for monitoring the health of physical point-to-point (port-to-port) network segment, but it does not provide visibility across the entire Ethernet service (EVC) from one Subscriber location to another.

For more information on 802.3ah Link OAM can be used as a troubleshooting tool, visit www.omnitron-systems.com and download the 802.3ah OAM white paper.

End-to-End Service OAM

The IEEE, the ITU and the Metro Ethernet Forum (MEF) have developed OAM standards to make Ethernet business services truly carrier grade. These standards bring the reliability of TDM and Frame Relay services to Ethernet Services, and deliver enforceable Service Level Agreements (SLA).

These standards have enabled End-to-End Service OAM.

For the Service Provider to have complete visibility across the entire network, including any portion of the network that traverses through an Out-of-Franchise region, the Service Provider must have End-to-End Service OAM tools. End-to-End Service OAM provides fault detection and performance monitoring of the entire Ethernet Services (EVC) from the UNI at one Customer Premises to the UNI at the other Customer Premises.

End-to-End Service OAM supports these functions across one or more Operator networks:

- Provide proactive fault monitoring
- Provide an efficient mechanism for rapid fault isolation to minimize down-time
- Ensure Class of Service consistency
- Support Service Level Agreement enforcement with meaningful reporting

The Goals of Service OAM are:

- Reduce operating costs associated with lengthy turn up of services and fault isolation
- Increase revenue with Service Level Agreement assurance
- Improve service reliability and Subscriber experience

Fundamentals of Service OAM

The two predominant OAM standards that have evolved to meet these needs are "IEEE 802.1ag Connectivity Fault Management" (FM or CFM), and "ITU-T Y.1731 OAM Functions and Mechanisms for Ethernet-based Networks", commonly referred to as Performance Monitoring (PM).

802.1ag provides connectivity checking and troubleshooting tools across multiple networks and multiple domains. Y.1731 expands upon 802.1ag concepts, and adds performance measurement and monitoring, giving Service Providers the tools needed for SLA assurance.

802.1ag and Y.1731 utilize OAM Protocol Data Units (OAMPDUs). OAMPDUs are artificially generated test frames (synthetic traffic), which follow the service path to detect faults and measure performance in real time. They are non-intrusive and non-traffic interrupting. 802.1ag and Y.1731 each have a different set of OAMPDUs, and each OAMPDU performs a different OAM function.

Multi-Domain OAM Model

One of the important benefits of 802.1ag and Y.1731 is their ability to support multiple domain levels. A domain level can be thought of as a viewpoint, or perspective, of different stakeholders in an Ethernet service (typically the Subscriber, the Service Provider and the Out of Franchise Operator.)

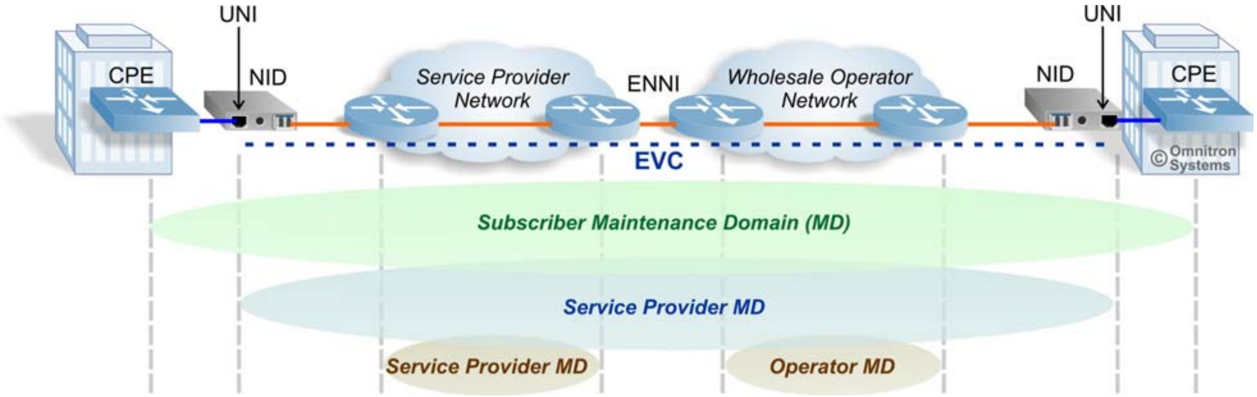

Figure 2. Multi-Domain Model

Maintenance Domains are abstract concepts that break the maintenance of Ethernet services into different levels. These levels make it easy to delineate the responsibilities and perspectives of the different stakeholders. Figure 2 illustrates three such domain levels, however, 802.1ag and Y.1731 can support up to eight levels of domains. The domains are numbered 0-7 and are assigned as needed to each stakeholder.

The top-level domain represents the Subscriber's perspective, shown in the green portion of Figure 2. The Subscriber domain traverses the entire length of the service from the Customer Equipment at one location to Customer Equipment at the other location. The Subscriber seeks assurance that the equipment is properly communicating from one side of the network to the other across the contracted service. The Subscriber's concern is that the service is operational, and that it performs at the contracted level. In case of service interruption or quality deterioration this tool enables the Subscriber to document and demand service quality improvement and potentially receive economic reimbursement.

The domain under the Subscriber represents the Service Provider's perspective (shown in blue), and spans from the UNI on the left to the UNI on the right. The Service Provider is focused on monitoring the availability and performance of the service from UNI to UNI to ensure the contracted service quality or SLA is met. Since the Service Provider is the primary contract service holder with the subscriber, the Service Provider is motivated to ensure that the SLA is met and repair any service flaws within his or his wholesale Out of Franchise Operator partner as soon as possible to enable quality of service to the Subscriber and minimize exposure to SLA penalties.

The domain at the bottom represents the Out of Franchise Operator's perspective, shown in beige, is the portion of the network that is controlled by each individual Operator. Each Operator is transporting a segment of the Ethernet service, and seeks to ensure their network is performing as contracted to the Service Provider. In the above example, there are two maintenance domains at this level because the Service Provider is also an Operator with their own local network, and must monitor that domain as well.

Maintenance Associations and Maintenance Points

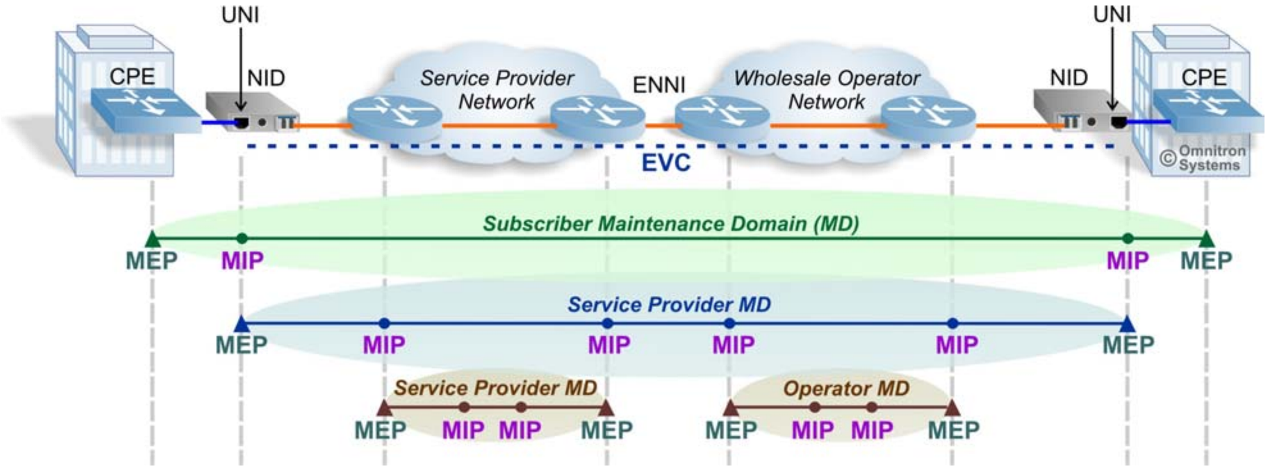

Figure 3. MEPs and MIPs

802.1ag and Y.1731 functions are based on the concepts of Maintenance Associations and Maintenance Points. Maintenance Associations (MA) are the physical network paths that reside in each domain. Here is where the different stakeholders define boundaries of networks and network links.

Any port on an MA can be designated as a Maintenance Point. Maintenance Points that reside at the outside edges of a MA are classified as Maintenance End Points (**MEP**).

When a Maintenance End Point is created, it is associated with a specific Maintenance Domain, Maintenance Association, Ethernet Service (EVC) and port number.

Maintenance Intermediate Points (**MIP**) reside between MEPs and provide points along the service path for isolating network problems. The MIPs in the blue Service domain represent edges of different Operator networks. Additional MIPs within the Operator networks may also be assigned. The MIPs in the Operator MAs at the bottom of the diagram represent Operator network equipment not illustrated within the Operator clouds.

These Maintenance Domains, Maintenance Associations and Maintenance Points define different network elements and enable the capabilities of End-to-End Service OAM.

IEEE 802.1ag Connectivity Fault Management

IEEE 802.1ag Connectivity Fault Management (CFM) is the foundation for Service OAM. 802.1ag is referred to as Connectivity Fault Management because it has the capability to proactively monitor and provide fault detection of the entire EVC. 802.1ag fault monitoring functions across all domains and Maintenance Associations. The Service Provider can use 802.1ag to monitor for faults from UNI to UNI, while each Operator along the service path can monitor for faults across their service responsibility from ENNI to ENNI or from an ENNI to a UNI.

802.1ag contains three important mechanisms:

- Continuity Check Message (continuous monitoring)
- Loopback (provider-initiated fault detection)
- Link Trace (provider-initiated fault isolation)

Continuity Check Message

A Continuity Check Message (CCM) is an OAMPDU that provides service monitoring from one endpoint 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.

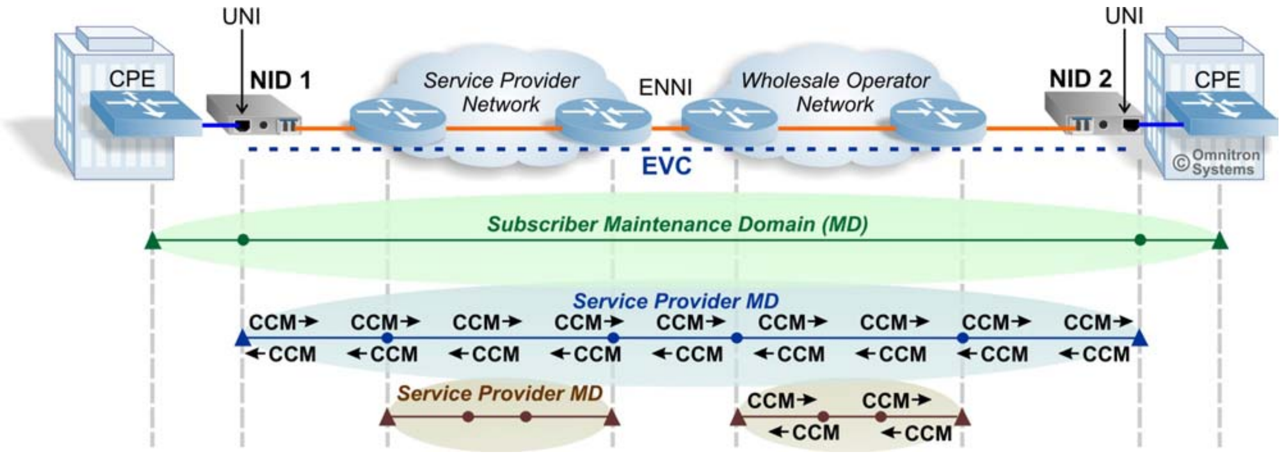

Figure 4. Continuity Check Messages

In Figure 4, CCMs are shown in two example domains. In the blue Service Provider domain, the Subscriber facing port on NID 1 (on the left) is configured as a MEP. This MEP periodically sends CCM frames to NID 2 (on the right), which is also configured as a MEP. The CCM mechanism provides a means to proactively and constantly monitor the EVC for connectivity failures between MEPs. As long as NID 2 is receiving a consistent flow of CCMs from NID 1, the network manager knows that the EVC has end-to-end continuity. If NID 2 stops receiving CCMs, it can send an alarm notification to alert the Network Management System that the EVC connection to NID 1 has been lost. CCMs are also sent in the opposite direction from NID 2 to NID 1 so the EVC service connectivity is monitored in both directions through the network.

CCMs can operate at the different domain levels, and are running on the Wholesale Operator's Maintenance Association on the bottom right of Figure 4.

Figure 5 illustrates a fault in the network. The fault is represented by the jagged yellow line. In this example, it is a uni-directional, or one way fault, and the location is unknown. Once NID 2 on the right has not received three consecutive CCMs it sends an alarm to the network manager. NID 2 continues sending out CCMs marking them with a Remote Defect Indicator (RDI) flag, notifying the MEP receiving the CCMs that there is a loss of service and prompting it to send a notification to its Network Management System.

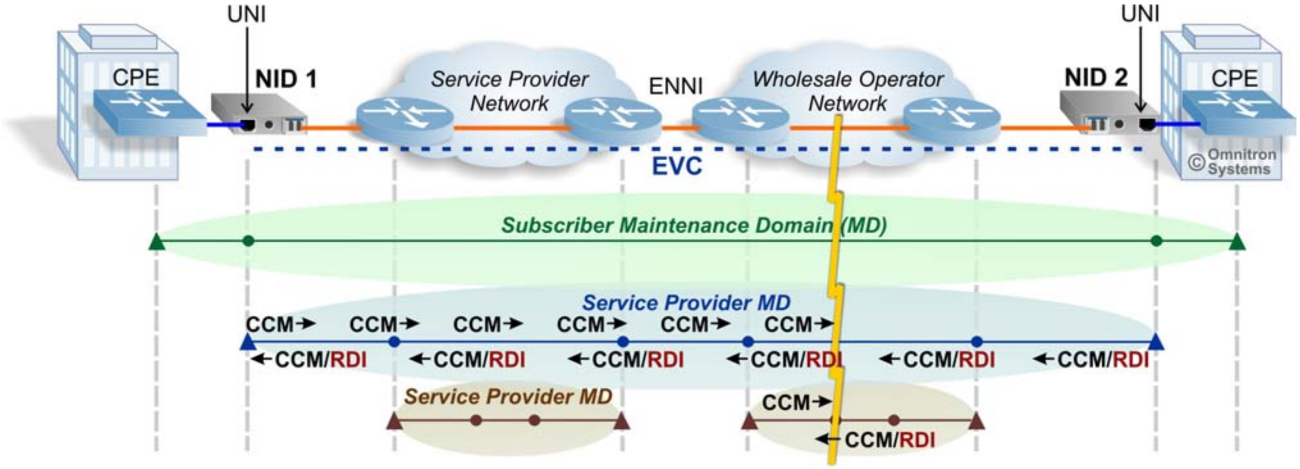
UNI
UNI
CPE
NID 1
Service Provider Network
ENNI
Wholesale Operator Network
NID 2
CPE
© Omnitron Systems
EVC
Subscriber Maintenance Domain (MD)
Service Provider MD
CCM→ CCM→ CCM→ CCM→ CCM→ CCM→
←CCM/RDI ←CCM/RDI ←CCM/RDI ←CCM/RDI ←CCM/RDI ←CCM/RDI
Service Provider MD
CCM→
←CCM/RDI

Figure 5. Continuity Check Messages with Network Fault

The benefit of Connectivity Check Messages is that network faults are proactively detected with an automatic mechanism. This enables instant notification of loss of service and gives Service Providers a head start verifying and restoring service as soon as possible.

Loopback/Fault Verification

Once it has been established that there is a network fault, the Service Provider may choose to verify the loss of service by initiating an 802.1ag Loopback test. With Loopback, a MEP sends messages to another MEP or MIP to verify connectivity across a given MA. Loopback is similar to a Layer 3 ping request/reply.

The mechanics of a Loopback are as follows: a MEP is selected and instructed to send Loopback Messages (LBM) to another MEP or MIP within the same MA. The MEP or MIP receiving an LBM returns a subsequent Loopback Response (LBR). LBMs/LBRs are used to verify bidirectional connectivity.

LBMs are typically initiated by the Service Provider, and can be provisioned to be sent as a one-time test, or to be sent periodically over time.

In some implementations, the 802.1ag Loopback function is used to provide two way (round trip) performance monitoring capabilities. It can measure delay, delay variation and frame loss of the LBM/LBR.

Link Trace

Once the Service Provider has confirmed there is a network fault, the Link Trace tool can be used to isolate the specific location of the fault. Link Trace is used to trace the service from one MEP to another MEP or MIP by its MAC address, and to all the MIPs along the MA.

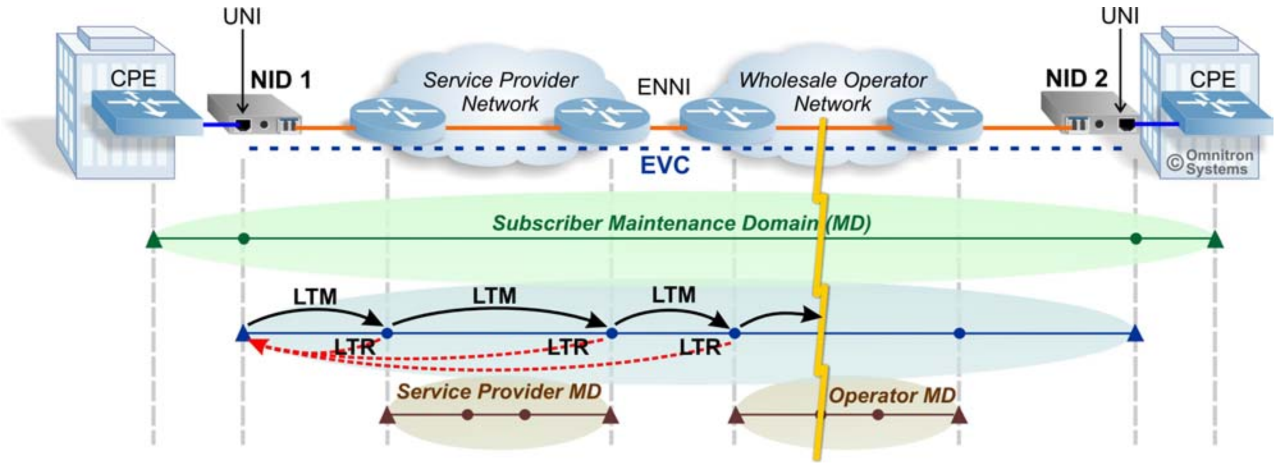

Figure 6. Link Trace

In the example above a network manager manually initiates a multicast Link Trace Message (LTM) from one MEP (NID 1) across the Maintenance Association. The message body of the LTM includes the destination MAC address of NID 2, which is the target MEP terminating the Link Trace. At each hop, when a MIP receives the LTM, it generates a unicast Link Trace Response (LTR) to NID 1 and forwards the LTM down the service path to the destination MAC address until the target MEP is reached. When the destination MEP receives the LTM, it too responds with the LTR back to the sender notifying it about the service path being functional. An LTM effectively traces the path from one MEP to the target MEP and allows the NID to discover vital connectivity data about the path and isolate the fault location.

In this example, the last MIP to respond with an LTM is at the edge of the Wholesale Operator's network at the ENNI, so the network manager can isolate the location of the fault to the Wholesale network. In some cases, when NID 2 is accessible via an out-of-band management channel, a Link Trace can be initiated from NID 2 to further isolate the fault.

The Wholesale Operator can also initiate a Link Trace from the MEP at the edge of his MA to isolate the fault within his network.

The Link Trace tool can also be used to establish a physical network path during service turn up by identifying relationships between remote MEPs and MIPs at the same domain level.

Link Trace provides the ability to quickly isolate the problem from any remote management location without deploying costly truck rolls. This capability improves Subscriber satisfaction with reduced down time and reduces loss of revenue from SLA enforcement.

ITU-T Y.1731 Performance Monitoring

ITU-T Y.1731 expands upon the concepts of 802.1ag and adds performance measurement and monitoring, giving Service Providers the tools needed for SLA guarantees and assurance. Like 802.1ag, Y.1731 can also operate at any of the eight administrative levels enabling the Service Provider and wholesale Operator(s), to monitor performance of the entire EVC, as shown in Figure 7.

Y.1731 provides the measurement of the following performance parameters:

- Frame Loss Measurement
- Delay Measurement
- Delay Variation Measurement
- Service Availability

Frame loss is the difference between the numbers of frames transmitted by one MEP at one end of the EVC and the number of frames received by the MEP at the other end of the EVC. Frame Loss Ratio is the ratio of frames lost to total frames sent.

Frame loss is calculated between MEPs by exchanging frame counters through Loss Measurement Messages (LMM) and Loss Measurement Responses (LMR). The MEPs compare frame counters and determine the number of frames lost.

Service availability can be determined by examining the Frame Loss Ratio over a period of time. Many Service Providers use Frame Loss Ratio to determine availability. If the Frame Lass Ratio exceeds a threshold, the service is considered unavailable.

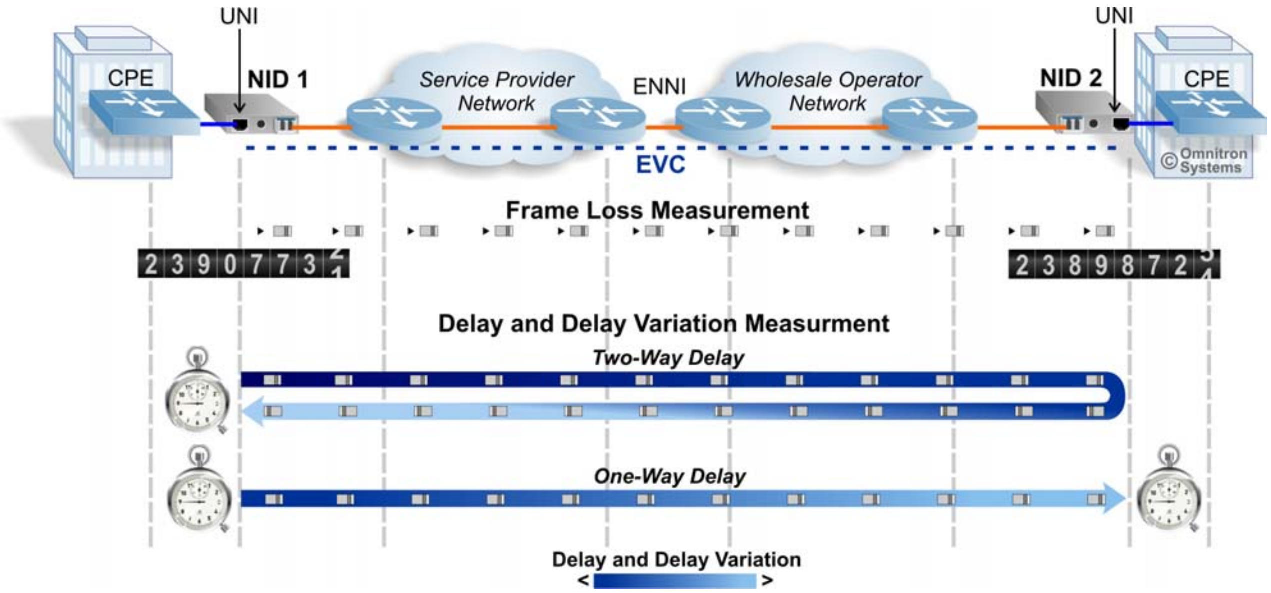

Figure 7. Frame Loss, Delay and Delay Variation Measurement

Frame delay is the time to deliver a frame from the source MEP to the destination MEP. There are two types of Frame Delay measurements, One-way and Two-way.

One-way delay is measured from the arrival of the first bit at the ingress MEP to the output of the last bit of the egress MEP. The source MEP will periodically send One-way Delay Measurement (1DM) frames, which include Transmit Time Stamps. The delay measurement is calculated at the receiving MEP by calculating the difference between the Transmit Time

Stamp and the Receive Time Stamp that is created when the 1DM frame is received. It should be noted that the One-Way-Delay measurement requires accurately synchronized time-of-day clocks between the two MEPs, which adds the expense of more equipment, and is less popular than the two-way delay.

Two-way delay does not require synchronized clocks because it measures round trip delay. It is accomplished by MEPs exchanging Delay Measurement Message (DMM) and Delay Measurement Response (DMR) frames. Each of these OAMPDUs include Transmit Time Stamps. Y.1731 allows an option to include additional time stamps such as a Receive Time Stamp and a return Transmit Time Stamp. These additional time stamps compensate for DMR processing time.

Frame Delay Variation is defined as the variation in delay for two delay measurements. Having a consistent delay (i.e. minimal Frame Delay Variation) is important for mobile backhaul applications and business services such as voice and video.

Some NIDs can capture and collect the data from performance monitoring and present it periodically to the service provider’s network management system. This data can be processed and presented in numerical or graphical form to the wholesale operator partner for SLA enforcement, or the Subscriber for proof of SLA conformance.

Comparison, Contrast and Clarification

802.1ag and ITU-T Y.1731 Functionality

- 802.1ag can manage faults in E-LAN services (multi-point to multi-point) whereas Y.1731 Performance Monitoring is limited to E-Line (point-to-point) services.
- 802.1ag can manage faults on an EVC basis. Per the ITU specification, Y.1731 can monitor performance metrics on an EVC basis, and it can monitor each Class of Service (CoS) within an EVC.
- 802.1ag Loopback may be used for two-way performance monitoring capabilities, but is user-initiated, and does not run continuously like Y.1731 Performance Monitoring.

OAM Function		802.1ag	Y.1731	Mechanism
CFM	Fault Detection	√	√	CCM
	Fault Notification	√	√	RDI
	Loopback	√	√	LBM/LBR
	Fault Isolation	√	√	LTM/LTR
	MIP Discovery	√	√	LTM/LTR Multicast
PM	Frame Loss		√	CCM, LMM/LMR
	Frame Delay		√	1DM, DMM/DMR
	Frame Variation		√	1DM, DMM/DMR

IEEE and ITU Terminology

The ITU and IEEE have different terminology for the same concepts, and in some cases different terms are presented with the same acronyms.

IEEE 802.1ag	ITU-T Y.1731
Maintenance Domain (MD)	No such term present
Maintenance Association (MA)	Maintenance Entity Group (MEG)
Maintenance End Point (MEP)	Maintenance Entity Group End Point (MEP)
Maintenance Intermediate Point (MIP)	Maintenance Entity Group Intermediate Point (MIP)

Service OAM and the Metro Ethernet Forum

While the ITU and the IEEE have developed a comprehensive set of standards that define the functions and mechanisms of Service OAM, the Metro Ethernet Forum (MEF) has and is developing Technical Documents that provide guidelines and recommendations on how to use these standards.

These include a Service OAM Performance Monitoring Implementation Agreement that will provide implementation constructs for Performance Monitoring, including specific recommendations on how to collect performance data over time and provide performance metrics and SLA reports.

The MEF is also developing an Abstract Test Suite to specify requirements and framework for Service OAM that represents expectations of Service Providers and Subscribers in managing Ethernet Services.

Network Interface Devices

The Network Interface Devices (NIDs), also known as an Ethernet Demarcation Devices, are normally placed at the Customer's Premises and provide the physical User-to-Network Interface (UNI) that delivers the Ethernet Service. They are normally owned and managed by the Service Provider, and used to provide the demarcation point where the service is handed off to the Subscriber.

NIDs support a variety of port configurations with one or more Subscriber ports (UNI) and Service Provider network ports. The physical interfaces on these ports can be RJ-45, fixed fiber or pluggable transceivers and support 10Mbps up to Gigabit data rates.

Figure 8. Omnitron's iConverter NIDs

End-to-End is UNI-to-UNI

Service OAM provides fault management and performance monitoring of the entire EVC from UNI to UNI, so the demarcation device at the Customer Premises must be service-aware. Service OAM functionality requires intelligent demarcation that supports the tools provided by 802.1ag and Y.1731. Demarcation NIDs are typically configured as Maintenance End Points (MEPs) to enable Service OAM functions.

NIDs have several important functions. They enable Service Attribute functions, such as EVC mapping, COS prioritization/mapping and Rate Limiting. NIDs may also provide protection options. But most importantly, NIDs provide the intelligent demarcation at the Customer Premises that enables End-to-End Service OAM.

Summary

This paper has provided an overview of how 802.1ag Connectivity Fault Management enables proactive fault detection and rapid fault isolation, and how Y.1731 Service OAM collects and processes performance monitoring statistics for meaningful SLA assurance.

The ITU, the IEEE and the MEF have played critical roles in developing these ground-breaking standards that have made Ethernet truly carrier grade, and as reliable as traditional TDM services.

The Goals of Service OAM are to reduce costs, provide SLA assurances and improve service reliability for a quality Subscriber experience. Service OAM standards achieve these goals with a robust set of functions that:

- Facilitate quick service turn up by helping provision the EVC and validating the service
- Monitor the service and report performance metrics and availability
- Automatically report service faults and service level deterioration and provide rapid fault isolation

Carrier Ethernet Service Providers want to offer a dependable and consistent quality of experience, whether the Ethernet service spans one or one-hundred Operator networks. They want to do this reliably and at low cost. Therefore, they need the end-to-end Service OAM tools for rapid fault detection and isolation, and also for performance monitoring and Service Level Agreement assurance.

About Omnitron Systems

Based in Irvine, California, Omnitron Systems designs and manufactures carrier-grade fiber connectivity products for Telecom, Enterprise and Government networks. Since 1992, network operators have relied on Omnitron products to enable the delivery of new services, and increase the capabilities of their fiber network infrastructure.

Omnitron's iConverter multi-service platform of Network Interface Devices, CWDM Multiplexers and T1/E1 Multiplexers support advanced fault detection and performance monitoring capabilities improve reliability and reduce operating costs.

The iConverter multi-service platform is used by Telecom Service Providers (ILECs and CLECs), Cable MSOs and ISPs to:

- Deliver Carrier Ethernet with end-to-end Service OAM
- Increase capacity of access, SONET/SDH and metro fiber networks
- Migrate from TDM to Ethernet/LTE mobile backhaul
- Extend demarcations in buildings and business complexes
- Extend distances in SONET/SDH ring networks
- Seamlessly integrate copper and fiber networks

MEF 9, MEF 14 and MEF 21 Certified Compliant

NEBS Level 3 Certified

Lifetime Warranty

Free 24/7/365 Technical Support

Made in the USA

CARRIER
ETHERNET
MEF
Certified Compliant

Omnitron Systems
Technology, Inc.
140 Technology Drive
Irvine, CA 92618 USA
www.omnitron-systems.com
info@omnitron-systems.com
800-675-8410
+1 949-250-6510