

Ethernet Aggregation and Transport Infrastructure OAM and Protection Issues

Pasula Reddy and Sam Lisle, Fujitsu Network Communications

ABSTRACT

As Ethernet traffic increases rapidly in carrier networks, service providers have an increasing need for Ethernet infrastructure networking in order to scale their Ethernet services. Ethernet infrastructure transports Ethernet traffic over distance, protects the traffic from link and nodal failures, and aggregates Ethernet connections from a large number of low-speed ports onto a much smaller number of high-speed ports. An Ethernet infrastructure allows service providers to inexpensively interconnect end user locations with one another for private line services, and to interconnect end users with the VPLS and IP/MPLS service edges for enterprise, residential, and mobility services. OAM and protection capabilities are critical to enable Ethernet to be deployed as an infrastructure technology. This article reviews several important developments in Ethernet OAM and protection standards, and discusses how those capabilities are vital for the creation of an effective Ethernet infrastructure.

INTRODUCTION

Ethernet traffic is entering North American provider networks at unprecedented rates. For example, retail enterprise Ethernet ports are projected to grow at a 40 percent compound annual growth rate between 2007 and 2012 [1]. Introduced initially in provider networks as a disruptive metro enterprise service, Ethernet is now expanding to inter-metro virtual private network (VPN) services as a layer 2 alternative to multi-protocol label switching (MPLS) VPN services, and has become the overwhelmingly dominant backhaul interface for residential broadband and triple play services delivery. Ethernet is also poised to become the dominant backhaul interface for mobile services with the pending deployment of fourth-generation (4G) technology.

As the volume of Ethernet traffic has increased and the applications for Ethernet have broadened, service provider Ethernet architectures have evolved. For example, initial provider Ethernet networks that provided intra-metro enterprise switched services consisted largely of Ethernet switching platforms in the metro core

running 802.1ad provider bridging protocol (or switch/routers running both Ethernet bridging and MPLS protocols), accompanied by simple media conversion or demarcation devices at the customer location. These bridged/routed switched services networks existed separately from networks providing layer 1 Ethernet port transport services that utilized synchronous optical network (SONET) or wavelength-division multiplexing (WDM) point-to-point connections.

These early switched services networks typically utilized no underlying transport among the switch/router service elements or between the switches/routers and the demarcation devices, so every customer port appeared as a port on the service element, and all interconnections between elements used dedicated optical fiber runs and were unprotected at the physical layer. As these networks grew, service providers deployed Ethernet port transport infrastructure that provided layer 1 protection and enabled multiple Ethernet ports among the switches or between customer locations and the service edge to interconnect over the same fiber. However, this transport infrastructure provided no Ethernet aggregation; therefore, each port at the customer location still appeared as a port on the switch/router service element. Similarly, as enterprises began to use Ethernet for Internet access and IP business services access, the transport infrastructure would typically carry each customer port to a dedicated port on the service edge element.

As service providers deployed residential broadband services using platforms that utilized Ethernet as the network interface, it became absolutely critical to not only transport Ethernet ports effectively, but to provide significant Ethernet aggregation back toward the service edge (Fig. 1). Provider networks often complemented layer 1 transport by using separate switch/router service elements to provide the aggregation and protection functions as Ethernet traffic was backhauled to the service edge. While these switch/router platforms often provided the essential aggregation and protection functionality, the operations, administration, and maintenance (OAM) and protection protocols involved were often more suited to routed network functionality rather than classic aggregation network functionality; therefore, the

performance, scalability, and cost of the solution were sometimes compromised.

Hence, as the volume of Ethernet traffic grows, Ethernet extends into additional applications such as mobile backhaul, and the need for higher-quality Ethernet private line services emerges, the OAM and protection capabilities of Ethernet reviewed in this article will become more and more vital for smooth operation and scaling of these large networks. It is these capabilities, along with efficient integration into the WDM-based transport layer, that will allow service providers to deploy integrated Ethernet aggregation and transport infrastructures that provide private-line-equivalent E-LINE services and unidirectional multipoint services, and cost-effectively bring large volumes of Ethernet traffic from end customer locations to the switched/routed service edge for E-LAN and IP services delivery.

ETHERNET OAM TOOLS

As current and next-generation services migrate to Ethernet, it becomes imperative for Ethernet to support a wide variety of OAM tools that enable providers to capitalize on the simplicity and flexibility of Ethernet, while enabling providers to precisely manage large Ethernet infrastructures. Since many important services are currently being delivered over SONET infrastructure, the Ethernet OAM toolkit needs to provide at least comparable functionality. In addition, these tools must enable providers to offer more measurable, yet granular and stringent service level agreements (SLAs) to their customers.

OAM protocols typically comprise the following four components:

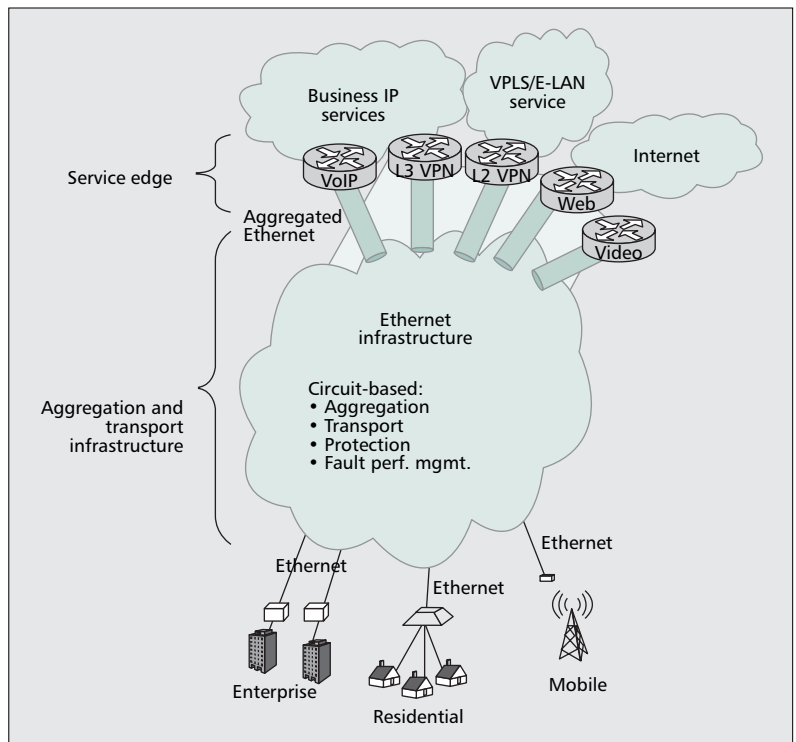
- Configuration and service provisioning
- Fault indication
- Diagnostic functions
- Performance monitoring

There are a number of OAM tools available for the Ethernet aggregation and transport infrastructure that can be used effectively to discover network elements, bring up and tear down services, monitor services individually, measure the performance against the SLA contract, and troubleshoot at the network, nodal, link, and per-service levels.

Work on packet-based OAM was largely pioneered for asynchronous transfer mode (ATM) technology, where a variety of mechanisms for fault and performance management, loopbacks, and other functions were developed [2]. Building on this conceptual foundation, the IEEE, International Telecommunication Union — Telecommunication Standardization Sector (ITU-T), and Metro Ethernet Forum (MEF) standards bodies have developed the following OAM protocols that can be used effectively in any network that uses an underlying Ethernet transport infrastructure to deliver services:

- IEEE 802.1AB [3]
- IEEE 802.3ah [4]
- IEEE 802.1ag [5]
- ITU-T Y.1731 [6]
- MEF 16 E-LMI [7]

These protocols enable provisioning, monitoring, and troubleshooting of E-LINE, E-LAN, and E-TREE services that are delivered completely within the Ethernet transport network. In



■ Figure 1. Ethernet aggregation and transport infrastructure.

Service		E-LMI	Y.1731
Network	802.1ag		
Transport		802.3ah	
Discovery	802.1AB		

■ Figure 2. Ethernet OAM protocols.


addition, they can also be used in Ethernet transport networks that are used to aggregate traffic from customer locations and hand it off to the service edge network elements (which are predominantly IP/MPLS-based) or backhaul mobile traffic from a base station at a cell tower to the mobile switching center.

These protocols operate at different layers within the Ethernet stack, as shown in Fig. 2, and serve different purposes, as discussed in the following paragraphs.

DISCOVERY LAYER

The discovery layer is not directly tied to the transport, network, and service layers. The protocols within the discovery layer assist in dynamically discovering attributes of physical links on network elements. This information is typically exported to the network management systems, and used for creation of topological maps and assisting in end-to-end path computation. IEEE 802.1AB (Link Layer Discovery Protocol) is used at this layer to discover physical links on network elements.

The transport layer is the physical or link layer within the Ethernet stack. IEEE 802.3ah Link Level OAM protocol is used at this layer for monitoring and isolating faults. In addition, IEEE 802.1ag can also be used at this layer for monitoring and fault-detection purposes.



■ Figure 3. 802.1ag operating over multiple operator domains.

TRANSPORT LAYER

The transport layer is the physical or link layer within the Ethernet stack. IEEE 802.3ah link-level OAM protocol is used at this layer for monitoring and isolating faults. In addition, IEEE 802.1ag can also be used at this layer for monitoring and fault-detection purposes.

IEEE 802.3ah operates on point-to-point links between Ethernet devices. Since it operates at the link or physical level, 802.3ah does not have any service awareness. 802.3ah relies on OAM protocol data units (PDUs) exchanged between the two Ethernet devices at either end of the point-to-point link. These OAM PDUs conform to the slow protocol exchange rates (maximum rate of 10 frames/s). As a result, the 802.3ah OAM PDUs can be generated and processed in software.

802.3ah supports the following functions:

- OAM discovery: Discover OAM capabilities on a peer device.
- Link monitoring: Event notification when error thresholds on the link exceed pre-set values.
- Remote failure indication: Notifies peer that the receive path is down or the link is slowly degrading in quality.
- Remote loopback: Puts the peer in intrusive loopback state to test the link and the peer. Statistics can also be collected while testing the link. These loopback messages are initiated on operator command.

The link events supported by 802.3ah include:

- Errored symbol period event: This event is triggered when the number of errored symbols exceed a preconfigured threshold within a window (measured in number of symbols).
- Errored frame event: This event is triggered when the number of errored frames exceed a preconfigured threshold within a time period (measured in 100 ms time intervals).
- Errored frame period event: This event is triggered when the number of errored frames exceed a preconfigured threshold within a window (measured in number of received frames).

- Errored frame seconds summary event: This event is triggered when the number of errored frame seconds exceed a preconfigured threshold within a window (measured in 100 ms time intervals).

Ethernet devices running the 802.3ah protocol can be in geographically disparate locations, enabling providers to monitor and isolate faults remotely without a truck roll.

NETWORK LAYER

The network layer deals with the forwarding of Ethernet frames based on tunnel identifiers within the frame such as VLAN tags. This layer could be used as the aggregation component for the service layer. For example, multiple EVPL services could be aggregated into the same Ethernet tunnel, where the tunnel is represented by the S-Tag in 802.1ad/point-to-point Q-in-Q networks or a B-Tag in 802.1Qay (PBB-TE) networks. (Individual EVPL service instances embedded within the tunnel are represented by the C-Tags in Q-in-Q networks or I-SIDs in 802.1Qay networks.) The IEEE 802.1ag connectivity fault management protocol is used at this tunnel layer for fault detection, network monitoring, and fault isolation.

The IEEE 802.1ag protocol enables providers to detect faults within milliseconds from the time they occur and also provides tools for isolating the faults. 802.1ag is a flexible hierarchical protocol that can be enabled at multiple levels and multiple layers [8]. It allows providers to partition their networks into multiple operational domains and end-to-end services to span multiple domains/carriers. For example, a service provider could provide an end-to-end EVPL service that spans two different operator networks, as shown in Fig. 3. 802.1ag allows each of the two operators to enable 802.1ag functionality independently within their networks, while also allowing the service provider and even the customer to enable end-to-end 802.1ag functionality that spans multiple operator networks.

Each node participating in an 802.1ag session is either a maintenance endpoint (MEP) or a mainte-

nance intermediate point (MIP). As the names suggest, these represent the ingress/egress nodes and transit nodes within a maintenance domain.

802.1ag protocol supports the following management functions:

- Continuity check messages (CCMs): These are exchanged among MEPs to detect loss of continuity or incorrect network connections. These messages contain Remote Defect Indication flags to report faults to other MEPs. CCMs can be sent every 3.3 ms, thereby ensuring that faults are detected within milliseconds from the time they occur.
- Loopback messages: These can be used to verify connectivity to remote MEPs and MIPs. Loopback messages are typically initiated by operator command as an in-service operation. These can also be used as an out-of-service diagnostic test.
- Linktrace messages: These messages are typically initiated by operator command and can be used to trace the path to remote MEPs and MIPs.

Providers can use 802.1ag management functions to monitor, detect, and isolate faults at the network layer. In Ethernet aggregation and transport networks, where the network layer could potentially be aggregating thousands of services into the same Ethernet tunnel, CCMs can be run for the tunnel at either 3.3 or 10 ms granularity rather than running CCMs on individual EVC services. This allows providers to monitor the Ethernet transport tunnels very aggressively and scales extremely well, since the total number of transport tunnels in an Ethernet aggregation and transport network is significantly lower than the number of services carried within those tunnels. In addition, these tunnel CCM messages can trigger protection switching immediately after the detection of faults, thereby allowing providers to meet 50 ms protection switching SLA requirements equivalent to SONET private line requirements.


SERVICE LAYER

The service layer deals with individual service instances, where a service instance represents a unique customer/subscriber and/or a unique flow within a customer/subscriber traffic stream. IEEE 802.1ag, MEF 16 E-LMI (Ethernet — Local Management Interface), and ITU-T Y.1731 are the protocols used at this layer for service provisioning, fault detection, and isolation within the scope of the service, service performance monitoring, and measurement.

E-LMI helps operators turn up services rapidly by automating the provisioning of Ethernet service attributes on attached customer premises equipment (CPE). It is an asymmetric protocol that allows the user-network interface (UNI)-N device to communicate relevant service related attributes to the CPE as shown in Fig. 4.

A few examples of some of these EVC attributes are:

- EVC state on the provider's Ethernet NE.
- UNI status: Conveys the UNI-N status and other service attributes of the UNI.
- C-VLAN ID to EVC mapping: This is used to convey information on how the CE-VLAN IDs are mapped to specific EVCs.



■ Figure 4. E-LMI operation.

- BW Profiles: Conveys bandwidth attributes such as CIR, CBS, EIR, EBS. These attributes can be used by the CPE to ensure that traffic originating from it conforms to the ingress bandwidth profiles agreed on in the SLAs.

IEEE 802.1ag can also be used at the service layer to monitor and troubleshoot individual service instances. Since protection switching is typically bound to the network layer, the CCM timers for the 802.1ag sessions running at the service layer are typically less aggressive, in the seconds/minutes interval granularity. 802.1ag sessions at the service layer can also be enabled on demand to monitor any connectivity problems associated with individual service instances that are carried within the Ethernet transport tunnels (which operate at the network layer).

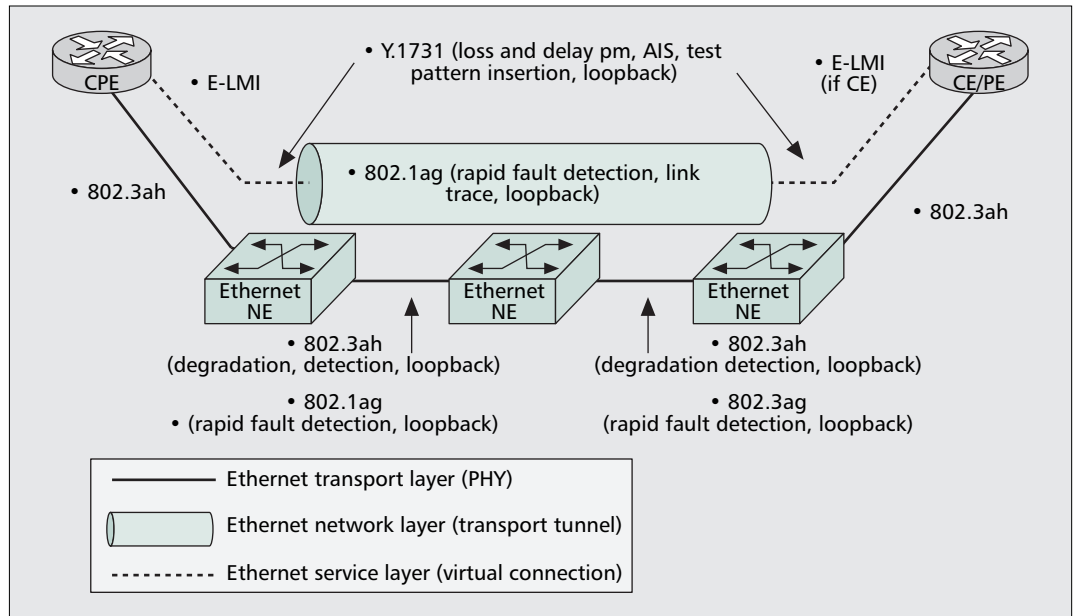
ITU-T Y.1731 also supports OAM functionality at the service layer. This protocol offers a few additional features on top of the IEEE 802.1ag protocol. It supports alarm indication signals (AISs), which can be used to propagate defect detection from a lower maintenance level to a higher maintenance level. When network elements receive AIS frames, the receiving network element records the AIS conditions, but does not generate loss of continuity alarms with peer MEPs at the local service layer. This benefits operators by suppressing potentially thousands of unwanted alarms on individual services that may be caused by underlying network or transport layer faults. Y.1731 supports measurements of performance parameters (frame loss ratio, one-/two-way frame delay, and frame delay variation) at the service instance granularity. These can serve as important tools for providers to measure their network performance and also document these measurements to prove to end customers that SLAs are being met. Typical implementations require specialized hardware assistance to measure the performance accurately. This can potentially translate to an increase in the cost of equipment that supports this functionality.

OAM PROTOCOL NETWORK APPLICATION

Having reviewed the major Ethernet OAM protocols, it is instructive to examine how these protocols work together in a simple example network environment. Figure 5 shows an Ethernet aggregation and transport infrastructure composed of Ethernet network elements bringing traffic from a customer edge (CE) router to another CE router or provider edge (PE) router.

E-LMI runs at the service layer over the Ethernet UNI between the CE router and the Eth-

E-LMI runs at the service layer over the Ethernet UNI between the CE router and the Ethernet network. This allows providers to turn up new services or modify existing services without dispatching a technician.



■ Figure 5. Ethernet OAM protocol application summary.

ernet network. This allows providers to turn up new services or modify existing services without dispatching a technician.

The 802.3ah link OAM protocol is used over all physical links in the network. This has particular value on UNI links to proactively monitor customer to network links and determine whether degradation defects exist within the provider network. 802.3ah can also be utilized on links within the network, but the slow protocol nature does not allow rapid detection of defects, so 802.1ag is also helpful on these links.

802.1ag operates on Ethernet transport tunnels at the network layer, enabling providers to monitor the Ethernet transport tunnel health, detect faults within the 10 ms benchmark established in SONET networks, and trigger protection switching onto a preprovisioned alternate path. Once faults have been detected (and traffic protected), the faults can be isolated using the loopback and linktrace messages.

802.1ag at the service layer monitors the connectivity at individual service instance granularity. The timers for the CC messages used at the service layer may be less aggressive than those for CC messages used in the network layer, because the network layer CC messages are used to trigger protection switching events on the transport tunnels.

Y.1731 can be used at the service layer to monitor connectivity, and also measure loss and delay performance at individual service instance granularity. Providers may choose to measure these PM characteristics for their high-touch customers with mission-critical applications.

Because the number of service instances can potentially run into hundreds of thousands within a large Ethernet aggregation and transport network, network providers need to consider the potential scalability implications if they enable very aggressive (millisecond granularity) 802.1ag or Y.1731 CC timers. In addition, if network providers do so, they also need to consider the

amount of OAM traffic traversing their network and take into account the bandwidth used by this OAM traffic in their call admission control (CAC) algorithms.

Occasionally, there is some confusion regarding the relative roles of IEEE 802.1ag and ITU-T Y.1731. Y.1731 comprises both fault management and performance management. The fault management functionality specified within Y.1731 is very similar to the IEEE 802.1ag specification, the only major exception being additional AIS functionality, which is present in Y.1731 but not in 802.1ag. Performance management functions such as frame delay measurements are only specified in Y.1731. Carriers looking to support performance measurement functions may be better served using Y.1731 for both fault management and performance management.

Depending on the service and application requirements, one or more of these tools can be used. For example, a service that offers E-LINE services completely within the Ethernet transport infrastructure network may enable the OAM protocols at all the layers, while a service that is backhauling broadband traffic to a BRAS may only enable the OAM protocols at the transport and network layers. The key point to note is that Ethernet provides OAM tools comparable to *non-packet-based* transport infrastructure OAM tools and also to other *non-Ethernet-based* packet transport infrastructure (e.g., ATM) OAM. These OAM protocols enable providers to implement a faster migration plan to an Ethernet-based packet transport infrastructure for their current and next-generation services/applications.

Interlayer OAM Relationships — Importantly, the OAM protocols at different layers complement each other. Defects detected within the OAM protocols at the lower layers can be propagated up to the higher-layer OAM protocols on an as-needed basis so that the higher-layer protocols can take appropriate actions. Some faults

such as link quality degradation may not necessarily be detected by higher-layer OAM protocols.

For example, 802.3ah link layer OAM at the transport layer may detect that the number of errored frames received within a time period has exceeded the configured threshold. This event needs to be propagated up to the network layer OAM protocol (802.1ag), which will enable it to set the Remote Defect Indication bit in the CCMs that it generates at the network layer. This allows the OAM protocol peers at the network layer to detect the fault. This information can be used as a protection switching trigger to switch the network layer transport tunnel to an alternative path, even before faults are explicitly detected using the triggers available natively at that layer. This could drastically minimize the service quality disruption times at the individual service instance granularity. It is important to note that, in this scenario, OAM protocols at the service layer will not declare any faults at that layer, preventing potentially thousands of unneeded alarms from being raised.

As discussed, 802.1ag CFM protocol can be instantiated at multiple layers. For example, 802.1ag can be enabled at the transport layer to detect loss of connectivity in tens of milliseconds to initiate a protection switching event. Although 802.3ah can also detect the loss of connectivity at the transport layer, there is a significant difference because 802.3ah link OAM is a slow protocol, which by definition implies that the failure detection times can only be on the order of hundreds of milliseconds at best. 802.1ag can also be used at the service layer to monitor and troubleshoot individual service instances, but it does not offer any PM capabilities. Y.1731, on the other hand, offers these PM capabilities, which can be used by providers to measure SLAs at service instance granularities.

Standards and Implementation Status —

Although the ongoing Ethernet OAM specification efforts within the IEEE and ITU organizations are progressing well, as of this writing, there is a need to define additional specifications focusing on OAM interworking functions between the Ethernet and IP/MPLS layers. As carriers migrate to Ethernet-based access and aggregation networks that feed into MPLS-based service networks, interworking between Ethernet OAM and MPLS OAM becomes very critical.

Since most services are typically individually identified across the inter-metro core at the MPLS layer and encapsulated in pseudowires, interworking between 802.1ag (or Y.1731 fault management functions) and VCCV is required to be able to propagate faults from one network domain to another, and also to provide end-to-end service traceability. If an MPLS network only relies on the Ethernet layer for providing layer 1/layer 2 connectivity between various MPLS-enabled network elements, the only required functions are monitoring and fault detection capabilities. 802.3ah and 802.1ag can be deployed to meet these requirements.

Concerning implementation status, as of this writing IEEE 802.3ah is being deployed by a large number of carriers for link monitoring functions. IEEE 802.1ag (or the fault management functionality within ITU-T Y.1731) is increasingly being

considered by various carriers as a mechanism to allow better service visibility to their customers. The performance management/measurement functionality specified within ITU-T Y.1731 is also being closely looked at by a large number of carriers, especially as they begin to offer new types of services that require meeting very stringent SLAs. Interoperability between different vendor implementations is an issue of concern to carriers, because they are forced to use the *lowest common denominator* capabilities in their multi-vendor networks. In some cases, this limits carriers from being able to measure and guarantee very granular SLAs. The same concerns apply to services that span multiple-carrier networks

ETHERNET PROTECTION PROTOCOLS

Service availability (uptime) and protection switching speed are critical requirements for Ethernet services and Ethernet-based access to IP services. As end users migrate mission-critical services away from TDM onto Ethernet, it becomes crucial for the Ethernet aggregation and transport infrastructure to offer availability and protection switching performance equivalent to that of SONET networks. Therefore, the Ethernet aggregation and transport network must be able to recover from link and/or nodal failures within 60 ms of a failure (10 ms for detection and 50 ms for switching). Ethernet spanning tree protocol (STP) variations are not capable of this level of performance.

There are two Ethernet forwarding mechanisms that provide Ethernet aggregation and transport infrastructure:

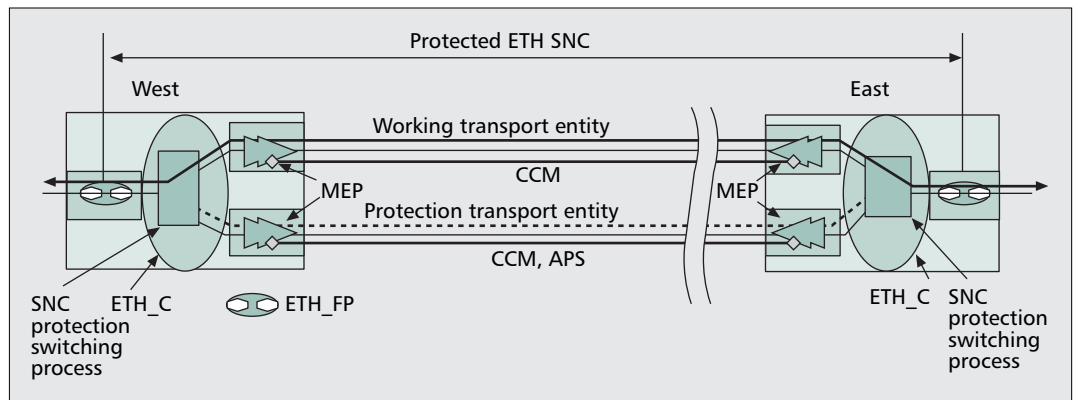
- Ethernet aggregation and transport over point-to-point VLAN(s)
 - Ethernet-based packet aggregation and transport over IEEE 802.1Qay (PBB-TE) [9]
- 802.1Qay is a more scalable technology, but service providers have a deployed base of VLAN-based aggregation and a desire to transition toward 802.1Qay over a period of time.

The point-to-point VLAN model aggregates and transports traffic from the customer edge to the service edge using VLAN tag(s). Customer traffic is isolated using a combination of VLAN tags, typically using Q-in-Q encapsulations. In the simplest case, services can be delivered over S-Tags. To increase the scalability, services can also be delivered over S- and C-Tagged connections. In this model both the S-Tag and C-Tag have significance within the provider network, enabling them to scale the number of service instances much more efficiently.

802.1Qay provides a more scalable forwarding mechanism that meets all the functional and availability requirements of Ethernet-based packet aggregation and transport networks. 802.1Qay uses an extended service identifier (I-SID) to embed individual services instances that exist at the service layer into a backbone tag (B-Tag) that exists at the network layer. It offers a highly scalable and efficient protection mechanism by ensuring that all the service instances between a given pair of ingress and egress nodes fully *fate-share* with the associated transport tunnel at the network layer. This allows the protection scheme to be completely implemented within the network layer, thereby reducing the number of protected sessions in the network drastically.

As end users migrate mission critical services away from TDM onto Ethernet it becomes crucial for the Ethernet aggregation and transport infrastructure to offer availability and protection switching performance equivalent to that of SONET networks.

The G.8031 protocol allows for 10ms fault detection and 50ms automatic protection switching that can operate at the service or network layers. It can be used at the network layer to provide protection for point to point VLAN based forwarding mechanisms and can optionally be applied to 802.1Qay networks.



■ Figure 6. G.8031 APS protocol.

Both forwarding mechanisms create point-to-point Ethernet networks with no MAC learning and flooding functions enabled. Therefore, protection mechanisms in these networks are not bound by xSTP reconvergence times while recovering from failures. Each forwarding mechanism has a protection switching capability that allows the transport infrastructure to deliver 50 ms automatic protection switching.

The G.8031 protocol [10] is a robust ITU standard that allows for 10 ms fault detection and 50 ms automatic protection switching that can operate at the service or network layers. This protocol can be used at the network layer to provide protection for point-to-point virtual LAN (VLAN)-based forwarding mechanisms and optionally be applied to 802.1Qay networks.

G.8031 supports the following protection modes:

- 1+1 bidirectional protection switching
- 1:1 bidirectional protection switching
- 1+1 unidirectional protection switching

G.8031 uses an APS PDU to signal to the far end that a protection switching event needs to be triggered. This is required to ensure that traffic is *co-routed* (carried along the same physical paths) in both the forward and reverse directions (for bidirectional service connectivity). Working and protect paths are preprovisioned, and both can be actively monitored using the CC messages (running at 3.3 or 10 ms granularity) in the 802.1ag CFM protocol. It is to be noted that the APS PDUs are only sent on the protect path as described in Fig. 6.

G.8031 supports a variety of triggers for protection switching. Some of the triggers are:

- Signal fail: This can either be detected locally or using 802.1ag CC messages.
- Signal degradation: This can be detected either locally or using any of the OAM protocols described in the previous section.
- Operator initiated switchovers: Operators can initiate switchovers manually to perform hardware upgrades and so on.

In addition, the G.8031 protocol supports both *revertive* and *non-revertive* modes of operation and options to control the time interval before a revertive switching event can be initiated.

The G.8031 APS function can be enabled at either the service instance granularity or on a group of service instances that fate-share the path between a given pair of ingress and egress

network elements, using the *test trail* functionality. Invoking the protocol at the network layer offers much higher scalability than invoking the protocol on a per service instance basis.

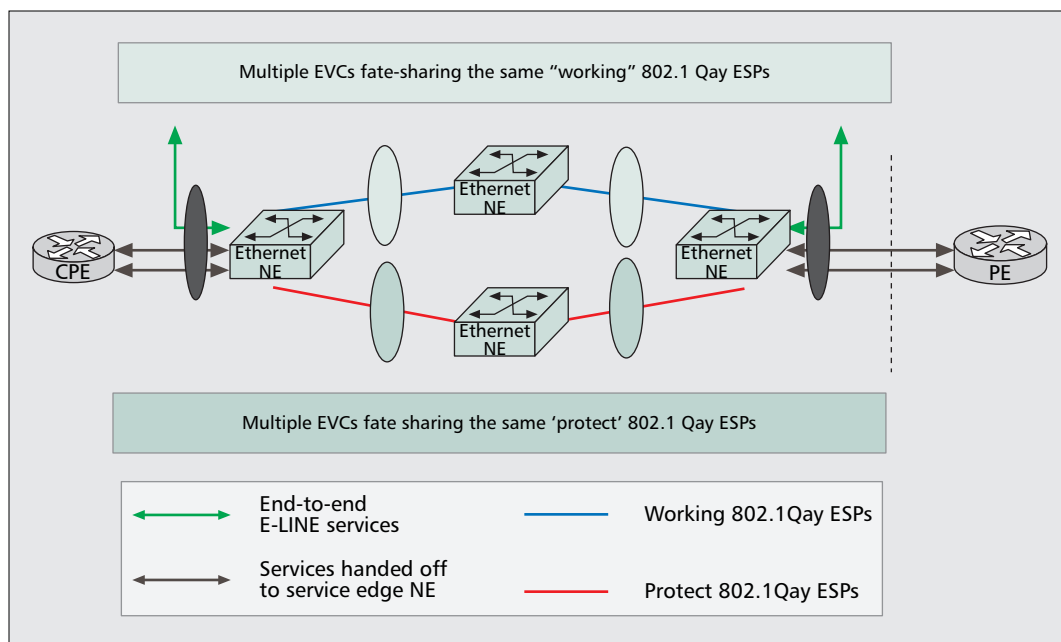
PBB-TE can optionally utilize G.8031, but also supports a simplified scheme that allows for load sharing. Working and protect paths are preprovisioned between pairs of ingress and egress network elements within the 802.1Qay network. Like G.8031, this mechanism also assumes that the paths in the forward and reverse directions are co-routed. The working and protect paths are monitored at the network layer using 802.1ag CC messages (running at 3.3 or 10 ms granularity). Once a failure is detected by the tail end node, it signals to the head-end using the RDI bit in the 802.1ag CC messages. This indicates to the head end that it needs to initiate a protection switching event. An APS PDU is not used in 802.1Qay protection switching. The additional flexibility 802.1Qay offers is load sharing support between the working and protect paths/tunnels. Like G.8031, 802.1Qay supports revertive and non-revertive modes of operation with options to control how long to wait before reverting to the working path. 802.1Qay supports only a 1:1 bidirectional protection switching mechanism.

Since failures are detected within tens of milliseconds from when they occur and the protection switching event is triggered as soon as the failures are detected, 802.1Qay networks recover from link and/or nodal failures within 50 ms. In addition, since the protection switching state machine runs at the network layer, it does not suffer from any of the potential scalability concerns from which other schemes that rely on protection switching state machines running at individual service instance granularity do.

As we can see, Ethernet offers two simple yet efficient protection protocols to meet aggregation and transport infrastructure requirements that could not be addressed by spanning tree protocols. This allows providers to migrate services from a SONET-based transport infrastructure to an Ethernet infrastructure without compromising protection switching speed.

CONCLUSIONS

As Ethernet applications continue to proliferate and traffic continues to grow, service providers must scale their deployments by constructing



■ Figure 7. Protection in 802.1Qay.

Of particular importance is the ability to operate OAM and protection protocols at the Ethernet network layer on aggregated tunnels where a small number of protocol instantiations provide fault detection and sectionalization and protection switching for a large number of Ethernet services.

Ethernet aggregation and transport infrastructure networks. These infrastructure networks cost-effectively aggregate, transport, and protect point-to-point Ethernet connections between end user locations providing native E-LINE services, and between end user locations and IP/MPLS/VPLS service edges. Just as SONET infrastructure has provided precision fault sectionalization, performance management, and rapid automatic protection switching, Ethernet infrastructure must provide similar functionality to support the many mission-critical applications of enterprise, wholesale, and mobility users in particular.

This article has reviewed key Ethernet OAM and protection switching protocols that have been standardized by the IEEE, ITU, and MEF. These protocols are essential enhancements to evolve Ethernet beyond a simple switched metro enterprise service. These protocols operate at various Ethernet layers, including transport, network, and service layers. The protection switching protocols enable dedicated 50 ms protection switching, which is identical in switching speed performance to SONET networks that have set the industry benchmark for protection performance. Of particular importance is the ability to operate OAM and protection protocols at the Ethernet network layer on aggregated tunnels where a small number of protocol instantiations provide fault detection and sectionalization and protection switching for a large number of Ethernet services.

REFERENCES

- [1] Vertical Systems Group, Inc. Emerging Network Services; retrieved May 2008, <http://www.verticalsystems.com>.
- [2] ITU-T I.610, "B-ISDN Operation and Maintenance Principles and Functions," Feb. 1999.
- [3] IEEE 802.1AB, "IEEE Standard for Local and Metropolitan Area Networks Station and Media Access Control Connectivity Discovery," 2005.
- [4] IEEE 802.3ah, "IEEE Standard for Information Technolo-

- gy — Telecommunications and Information Exchange Between Systems — Local and Metropolitan Area Networks — Specific Requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks," 2004.
- [5] IEEE 802.1ag, "Virtual Bridged Local Area Networks — Amendment 5: Connectivity Fault Management," Draft 4.1, Aug. 2005.
- [6] ITU-T Y.1731 "OAM Functions and Mechanisms for Ethernet Based Networks," Draft, ITU-T SG 13 WP 4 / Q5, June 2005.
- [7] MEF 16, "Ethernet Local Management Interface (E-LMI)," Jan. 2006.
- [8] D. O'Connor, "Ethernet Service OAM — Overview, Applications, Deployment, and Issues," *Tech. Dig. OCF/NFOEC*, Anaheim CA, Mar. 2006, p. 10.
- [9] IEEE 802.1Qay, "IEEE Standard for Local and Metropolitan Area Networks — Virtual Bridged Local Area Networks — Amendment: Provider Backbone Bridge Traffic Engineering" Draft 3.0, Apr. 18, 2008.
- [10] ITU-T G.8031 "Ethernet Protection Switching," June 2006.

BIOGRAPHIES

PASULA REDDY is responsible for data planning at Fujitsu Network Communications. Prior to joining Fujitsu, he worked in the product management organization at Redback Networks, where he focused on emerging layer 2 and layer 3 technologies. Before joining Redback, he worked at Torrent Networks, CoSine Communications, and the ProCurve Networking Business Unit at Hewlett-Packard. At ProCurve he concentrated on technologies related to high-end Ethernet switching and compliance with the appropriate IETF and IEEE standards. He holds a B.Tech. degree from the Indian Institute of Technology Madras and an M.Sc. degree from North Carolina State University.

SAM LISLE (sam.lisle@us.fujitsu.com) is a market development director at Fujitsu Network Communications, where he focuses on packet optical networking technology and applications. He holds U.S. patents in packet-aware optical networking and was instrumental in the definition of the Fujitsu FLASHWAVE® 4000 series of MSPP platforms. Prior to Fujitsu, he worked for Bellcore (now Telcordia), specializing in the reliability analysis of SONET transport systems and fiber optic media. He holds a B.Sc. degree in electrical engineering from the University of Iowa and an M.Sc. degree in electrical engineering from Georgia Tech.