Comparison of Ethernet and SONET OAM Status Report on Carrier-Class Ethernet OAM

Introduction

The telecommunications industry has a significant level of interest in Ethernet services, which include E-Line and E-LAN. These new services represent a potential vehicle for carriers to increase revenue. Some projections show a migration from SONET-based leased line services to Ethernet services, with Ethernet transport gradually replacing SONET transport. A debate exists in the industry with respect to how fast this trend will emerge, and what requisite requirements exist for Ethernet networks and equipment.

Enhancement of Ethernet OAM to become carrier class has generally been viewed as a key factor that will pace the growth of metro Ethernet services. The principal motivation for robust Ethernet OAM capability is to enable network operators to automate operations and administration, verify network performance, rapidly isolate problems and reduce operational costs.
Key goals include:

- Reducing operating costs by providing efficient means of failure/defect detection, diagnosis and handling
- Improving network and service availability
- Providing SLA performance monitoring and verification
- Ensuring security of customer traffic
- Minimizing customer problem reporting

To achieve these goals, a new OAM toolkit must be developed to support Ethernet services. The emerging importance of Ethernet services has spurred related standards forum activity in multiple venues. The SONET migration scenario has sparked discussion of whether Ethernet OAM is as robust as SONET OAM.

The goals for this paper are to summarize the current status of Ethernet OAM for service provider applications, compare Ethernet service OAM with SONET OAM and summarize conclusions. In order to accomplish these goals, the following methodology will be applied:

1. Define a service network model for point-to-point (E-Line) and multipoint (E-LAN) Ethernet services and summarize layering aspects.
2. Define the major OAM flows and functions that are required to support carrier-class Ethernet OAM.
3. For the major Ethernet OAM flows and functions, summarize the current standards and compare them to similar SONET flows and functions. Highlight standards status, existing functional gaps and unique Ethernet attributes.

Ethernet Service Network Models, OAM Flows and OAM Functions

Ethernet is both a LAN/MAN technology and a service. Ethernet technology can be used to implement Ethernet services, but connectionless Ethernet and associated bridging technology is not sufficiently scalable to build service networks, which span multiple cities and support many customers. MPLS-based connection oriented technology is required to scale Ethernet services. In addition, SONET transport will play a role in access networks. Figure 1 provides a reference model for Ethernet service networks.

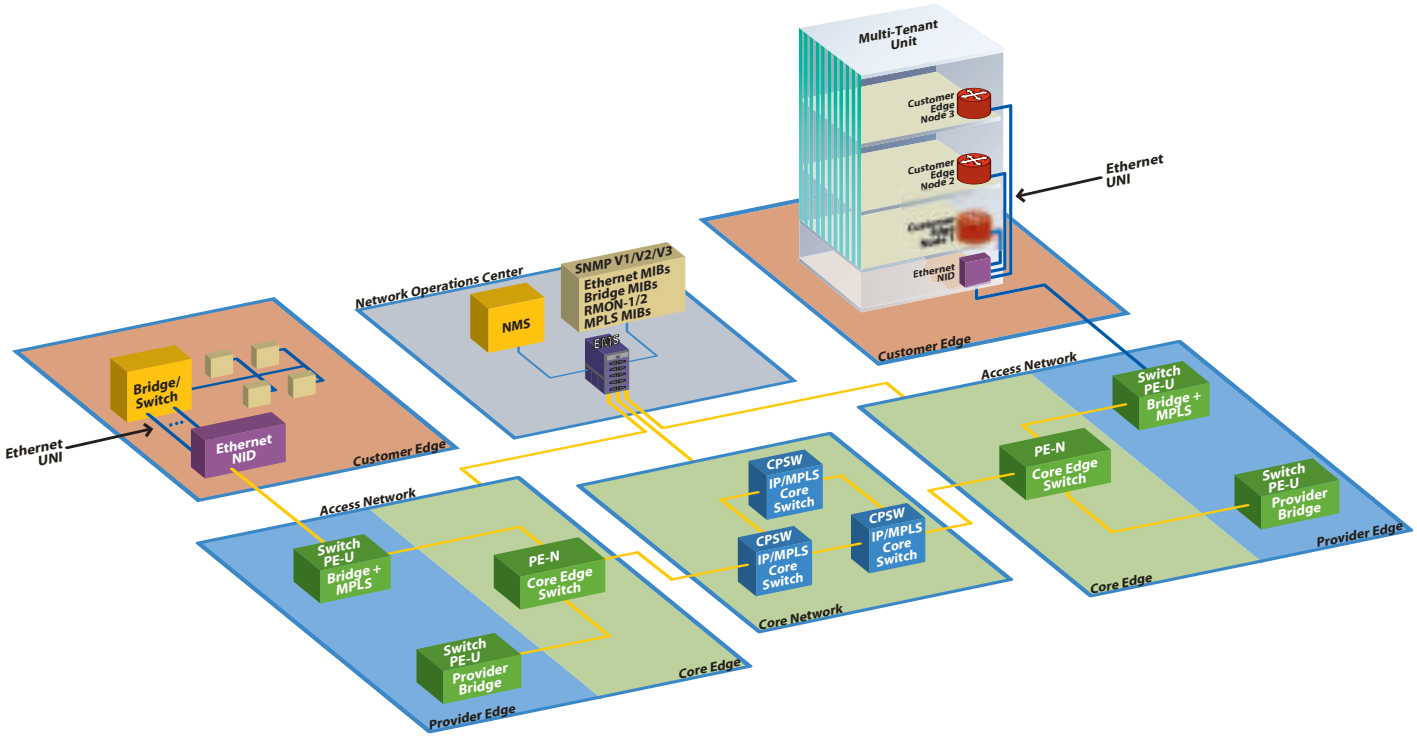

Figure 1: Ethernet Service Network Model

The network will support both E-Line and E-LAN service. The core network will generally be based on MPLS. Access networks will be based on multiple technologies, which include Ethernet provider bridges (as is currently being standardized by IEEE 802.1), MPLS and EoS.

The MPLS approach for Ethernet service transport uses a concept developed by the IETF known as pseudowires. A pseudowire emulates the attributes of a point-to-point service connection (Ethernet, frame relay, leased line) over a packet switched (IP/MPLS) network. Pseudowires are implemented by assigning an MPLS label to designate a virtual circuit encapsulated within a tunnel across a packet switched network. The tunnel is typically implemented by an MPLS label, although the tunnel can also be implemented by L2TP or IP-in-IP.

In order to support E-LAN service in a scalable manner, hierarchical/distributed hub and spoke architectures will be required in the access network. These architectures use a combination of bridging and MPLS switching. Relevant standards work is currently taking place in the IETF PPVPN group and the IEEE 802.1ad Provider Bridge group. The IETF and IEEE are working together to synchronize this work. The IEEE is defining the relevant bridging functions and the IETF is defining the relevant MPLS functions.

The IEEE architecture supports service networks where bridging plays a prominent role (i.e., larger bridge-based access networks). By contrast, the IETF hierarchical/distributed model divides access networks into PE-U nodes, which are customer facing and PE-N nodes, which are core network facing. PE-U nodes implement bridging and PE-N nodes may or may not implement bridging. Spoke connections between PE-U and PE-N nodes can be implemented by either MPLS pseudowires or an additional provider VLAN tag per the IEEE Provider Bridge Standard. This approach is known as Q/Q where an additional provider VLAN tag is stacked over an 802.1Q customer VLAN tag. SONET transport can be used to connect subscribers to the PE-U and to connect PE-U to PE-N. This layer is not shown in the diagrams for simplicity.

The MEF has developed a layered model for Ethernet service. This model is shown in Figure 2.

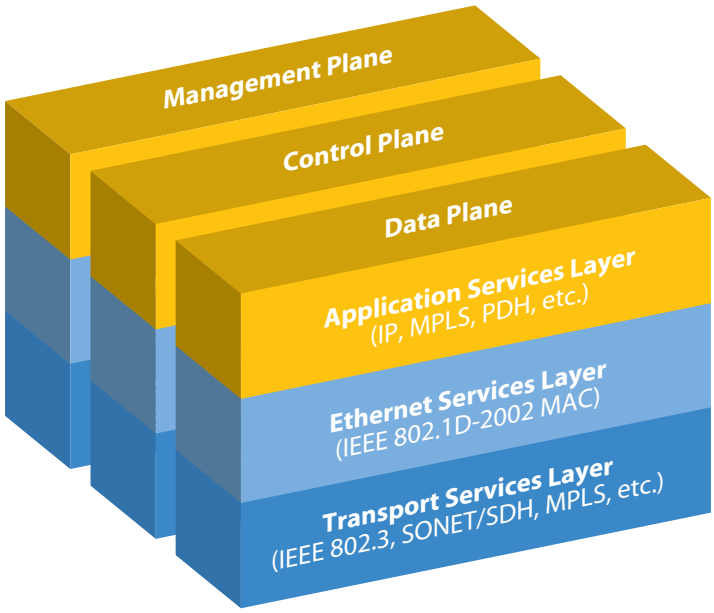

Figure 2: MEF Layered Model for Ethernet Service

The goal is to interconnect subscriber Ethernet ports and support application services such as IP. Ethernet services can be transported over either a packet switched transport layer, a TDM/SONET-based transport layer or a combination of the two. In order to administer and manage these service networks, OAM flows are needed at all three layers. This paper focuses on the requisite OAM flows at the Ethernet services layer and the transport services layer.

Figure 3 applies the layered service model to the Ethernet service network model and provides a recommended model for the set of Ethernet-related OAM flows required to fully support carrier-class Ethernet services. The NE-NE OAM flows must be defined for particular network segments and each layer of the model.

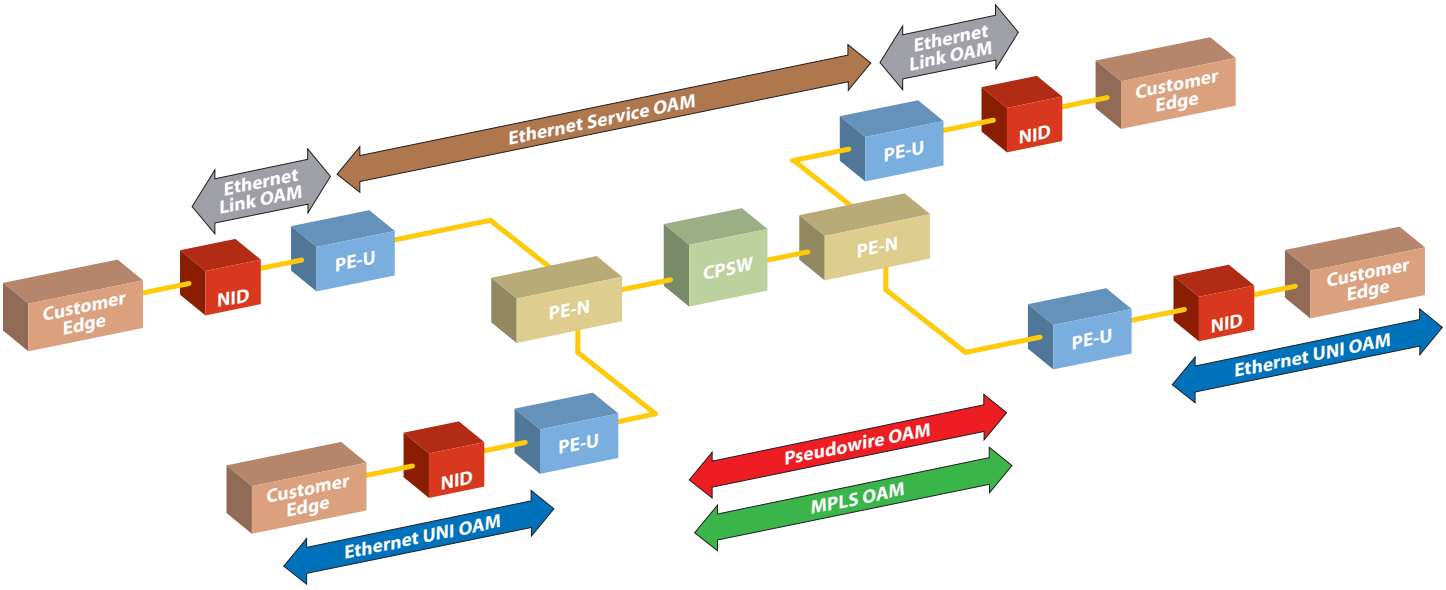

Figure 3: Segment and Layer OAM

In order to support carrier-class Ethernet service, the minimum core set of requisite OAM flows is:

1. Ethernet Service OAM PE-PE – needed to monitor and verify the integrity of Ethernet virtual connections from provider edge node to provider edge node.
2. Ethernet PE-NID Link OAM – needed to monitor and verify the integrity of the Ethernet link between the PE and the NID, which provides a managed demarcation between the provider network and the customer network.
3. Ethernet UNI OAM – needed to provide customers with service-related status information.
4. Transport Layer OAM – needed to monitor and verify the integrity of transport layer segments such as PE-U to PE-N links and PE-N to PE-N links.

With respect to the transport layer, this paper focuses on MPLS/pseudowire and provider bridge Q/Q based transport. For this transport layer, relevant OAM flows are MPLS tunnel OAM and pseudowire or Q/Q OAM.

For purposes of comparing Ethernet OAM to SONET OAM, we need to define an analogy between the layered/flow model (described earlier) and the SONET layered model. The physical and packet (SONET Path, GFP, PPP) transport layers are equivalent to the SONET section layer OAM. MPLS layer OAM is equivalent to SONET line layer OAM. Pseudowire layer OAM or provider bridge Q tag OAM is equivalent to SONET STS path layer OAM. EVC service layer OAM is equivalent to SONET VT path layer OAM.

With respect to the TMN functional areas, this paper focuses on fault management, performance management and security management. Configuration management is not addressed because control plane capability also plays an important role in this area, and the scope is limited to the management plane. Ethernet and MPLS technology include an extensive amount of control plane capability, which facilitates plug and play. However, standards have recently been generated for a GMPLS-based control plane for SONET, which offers the future promise of more automated service provisioning for these service networks. Accounting management is not considered because this area is primarily within the domain of higher layer OSSs.

For fault management, the following functions are considered:

- **Continuity Verification** – hello/status OAM messages are continuously exchanged between layer or segment end points
- **Ping/Loopback Test** – echo/reply messages or a path/segment loopback are used for fault isolation
- **Traceroute** – ability to discover the path for a service connection or a transport tunnel
- **Multi-Layer Alarms** – ability to detect network faults and generate forward and backward failure indications across multiple layers

For performance management, the following functions are considered:

- **SLA Verification** – accomplished by collecting service-related performance data
- **Network Monitoring** – gathering network performance data to provide an early an problem warning and assist with traffic engineering
- **Performance Alarms** – generation of threshold crossing alerts when performance data exceeds configured thresholds

For security management, the following functions are considered:

- **Link Authentication** – verify the integrity of a network attachment link
- **EMS/OSS to NE Security** – confidentiality and integrity verification for OSS to NE communications over an untrusted network

Ethernet Service OAM PE-PE Flow

Continuity Verification

The continuity of EVC between PE nodes must continuously be verified. This process is accomplished by exchanging hello OAM messages between EVC endpoints in ingress and egress PEs at some predetermined rate. The MEF is currently developing an IA for the Ethernet Service OAM flow [MEF Service OAM]. ITU SG 13 Question 3 Rapporteur Group has also started to work on Ethernet OAM. The two groups have initiated a liaison process to coordinate their efforts. At present, the MEF service OAM work is more advanced.

Hello frames are sent in-band and are differentiated from customer traffic by assigning an OAM Ethertype for the frame. Figure 4 illustrates the MEF Ethernet service OAM frame format.

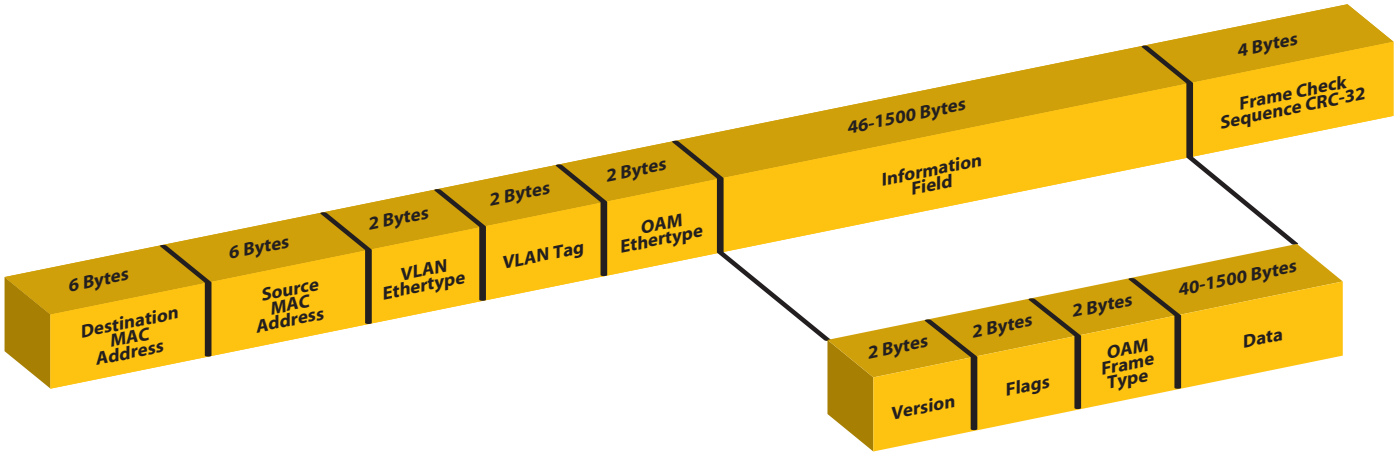

Figure 4: MEF Ethernet Service OAM Frame

The information field of the Ethernet service OAM frame indicates the type of OAM frame (e.g., hello), the version of the protocol, flags and any OAM frame type specific data.

The EVC to be tested must be distinguished. In order to conserve destination MAC addresses used for OAM, transport layer encapsulation is utilized. The hello message is sent along the same pseudowire or provider bridge Q tag as is used for customer traffic. EVC transport uses a combination of connection oriented pseudowires and connectionless provider bridges. In addition, EVCs can be either P2P or MP2MP. The net result is that a simple 1:1 mapping between a service OAM flow and an EVC does not exist. A given OAM flow tests multiple EVCs (for P2P) or EVC endpoints (for MP2MP) dependent on network architecture.

For some E-LAN service architectures (e.g., PE-Ns, which do not support bridging), one pseudowire per destination PE-U exists. For this case, each pseudowire service OAM flow tests the continuity of all the PP EVCs that terminate in that PE-U or all the MP EVC endpoints that terminate in that PE-U. For some E-LAN service architectures (e.g., PE-Ns, which support bridging), one pseudowire or provider Q tag per customer VPN exists.

For these cases, each pseudowire or provider Q tag OAM flow tests the continuity of all the PP EVCs or MP EVC endpoints for that VPN that are sourced by the PE-U which sources the OAM flow. The source and destination MAC address for the service OAM hello frame is a MAC address that is assigned to the PE node that sources or sinks the OAM frame. PE MAC addresses can potentially be learned if service OAM includes broadcast status messages. Such status messages would be broadcast to all other PEs in the network (using a well known multicast MAC address) and would include their unicast destination MAC address. Inclusion of such broadcast status messages is under discussion in the MEF.

For SONET service networks, service continuity is tested by using the J byte-based path trace function. The service path source NE can be configured to encode a path source ID in the appropriate path overhead J bytes. The service path sink NE can be configured to monitor the path J bytes for a particular source ID. If the correct path source ID is not detected, a path mismatch alarm will be generated. This process provides an effective service continuity check but is somewhat complex to provision. The principal difference is that the SONET approach has been standardized for some time, and the Ethernet approach is not yet complete (but when completed, the Ethernet approach will be simpler to administer).

Ping/Loopback Test

MEF service OAM also includes a ping test. A PE node sends a connectivity test request to another PE node, and that node returns a connectivity test response. The request/response messages are sent on the transport layer encapsulation to be tested. Multiple EVCs or EVC endpoints are tested per the discussion in the previous section. PE nodes can be addressed by either unicast or multicast virtual MAC addresses. If a unicast destination MAC is used, the connectivity test will be point-to-point. The two options for multipoint testing include: 1) use global multicast MAC address in the destination MAC field and per VPN multicast destination MAC in the OAM frame payload 2) use global multicast MAC in the destination MAC field and a unicast destination MAC in the OAM frame payload. The first option accomplishes a multipoint test across an entire VPN. Responses are generated after a random delay to prevent the initiating PE from being overwhelmed with test messages. For the second option, only the PE whose virtual MAC address matches the destination MAC in the test frame payload responds.

The connectivity test can also be used for performance management to measure delay and jitter. This process is accomplished by including a time stamp of when the frame was sent in the request payload and time stamps of when the request was received and the response sent in the response payload.

SONET path connectivity is tested by using a loopback function. GR-253 specifies five types of loopback tests: 1) SONET terminal loopback at the line layer 2) SONET facility loopback at the line layer 3) DSn terminal loopback 4) DSn facility loopback and 5) NE internal loopback test. For SONET terminal loopback, the received signal after O/E conversion and regeneration is looped back to the transmit path. For facility loopback, the received signal is looped back before regeneration. DSn terminal loopback takes place at the terminal side, and DSn facility loopback takes place at the line side. Network management signaling (rather than NE to NE OAM signaling) is used to put the SONET NE in loopback. The DCC cannot be used to carry the loopback commands for cases where the loopback would interrupt the DCC.

When comparing SONET loopback with Ethernet ping, Ethernet offers the following advantages: 1) Ethernet does not require an external test set to test the path 2) Ethernet can also measure delay and jitter 3) Ethernet is capable of multipoint testing. Delay is more of an issue with packet transport and SONET transport is PP.

The IETF PPVPN work group has also started some initial work on OAM for Ethernet L2 VPN services. IETF draft-stokes proposes a methodology for testing the hierarchical E-LAN architecture in which the PE-N includes bridging. Currently, this draft does not include a continuous hello continuity check but does include spoke-to-spoke ping and traceroute functions. A UDP OAM packet is sent end-to-end from spoke PE-U to core PE-N to far end core PE-N to far end spoke PE-U. Pseudowire transport and PE-U/PE-N data plane functions are tested end to end.

The IETF draft-stokes approach is a multi-layer OAM flow that is not solely at the Ethernet service layer as is the case with MEF service OAM. IETF draft-stokes extends MPLS layer OAM as developed by the IETF [IETF draft-mpls-lsp-ping]. IP/UDP packets instead of Ethernet frames (as is the case with MEF) are used for pings. Pseudowires, plus the bridge path, are tested simultaneously. As standards work progresses across multiple forums, concerns arise with violation of the principal of layering. At present, a clear divergence between the MEF approach and the IETF draft exists.

The UDP OAM packet is encapsulated in an OAM Ethernet packet whose destination MAC is either a MAC address assigned to the spoke or to a customer destination MAC that passes through that spoke. OAM packets are recognized at the far end PE-U by adding a special MPLS router alert label below the pseudowire label. The UDP OAM packet includes a field that indicates the reply mode for the ping. Typically, the reply will be sent over the spoke return path for the hierarchical E-LAN network. If this mode fails due to a fault on the return path, the reply can be forwarded in the non-MPLS IP data path or over the RSVP-TE control plane. The UDP OAM packet also includes time stamp fields so that delay and delay jitter can be measured like MEF service OAM.

Traceroute

Traceroute is an OAM function that has been successfully used in IP networks for many years. Ping packets are successively sent with incrementally longer TTL field values 1 to N where N is the length of the path in hops. TTL is decremented at each hop and when it equals zero, the packet is discarded and an ICMP control message is returned with the identity of the hop at which the packet was discarded. By using a succession of such pings, the router path can be traced to a destination IP address.

Traceroute is potentially useful as a fault management tool for Ethernet service networks. However, Ethernet frames do not have a TTL field, so the IP methodology cannot be directly adopted. The MEF has started to study the requirement, but the current plan does not include Ethernet traceroute capability in the first release of the EE service OAM implementation agreement. The challenge for Ethernet traceroute is to accomplish this function solely in the data plane, which involves the hardware-forwarding path and consequently involves potential changes to Ethernet switch chips. A control plane-based traceroute can be more readily accomplished with existing hardware but is not an optimal solution as the control plane's perception of the forwarding path may not always match the data plane reality.

An ideal Ethernet traceroute solution will be one that can be implemented in the control plane in the near term and in the hardware-based data plane in the longer term. The previously discussed MEF Ethernet service OAM frame can potentially be used for traceroute. Definition of payload content, which can be used to trace the route of the Ethernet service path, is the next step.

SONET OAM does not include a function that is equivalent to the traceroute. Path trace, which can be used for a continuity check, is not adequate because the path trace comparison is only made by nodes that terminate the path. The previously described loopback function is the closest equivalent for fault isolation tasks. Over the longer term, traceroute is one area where Ethernet can exceed SONET with respect to OAM capability.

IETF draft-stokes also include a spoke-to-spoke traceroute function. As with the ping function, this draft extends the label switched path traceroute function as defined in IETF draft-mpls-lsp-ping. A UDP echo packet with an Ethernet encapsulation is sent to a destination MAC address that corresponds to the far end spoke. This packet is sent to all adjacent nodes that are on the path to the far end spoke. The TTL field in the corresponding pseudowire label is set to 1. The packet will not be forwarded beyond the adjacent nodes because they will decrement the TTL to 0. The adjacent nodes return an echo reply packet whose payload contains the IDs (IP address) and associated label stack to reach the next nodes on the path to the spoke. The ingress spoke node then sends a successor echo packet with TTL = 2 in the pseudowire label and with the payload including the next hop information which was returned from the previous probe. The process continues until the far end spoke is reached. This series of pings will trace and validate the data plane spoke-to-spoke connectivity. The only issue is that this approach is somewhat orthogonal to the MEF layered approach and effectively mixes IP, Ethernet and MPLS layers.

Multi-layer Alarms

Multi-layer alarms have been one of the strengths of SONET OAM. When a fault occurs in the network, SONET NEs provide both forward and backward failure indications with linkage across all layers. The SONET alarm notification provides complete coverage and precludes cascading error conditions and alarm floods. This process is accomplished by propagating alarms between NEs, in addition to NE to OSS alarms, and linking alarms across layers within an NE. A condition detected at layer N causes alarms to be generated at layer N+1. NE-to-NE alarms are communicated by using overhead bytes and overhead byte content is the primary source of alarm triggers. NE-to-NE alarms are the focus of this section. NE-to-OSS alarms will be discussed in the performance management section.

The four primary areas of NE-to-NE alarms are: 1) Loss of signal/frame/pointer – remote defect indication, 2) Bit error related (as detected by overhead bit interleaved parity) – remote error indication, 3) Trace identifier mismatch (previously discussed in continuity section) 4) Unequipped or signal label mismatch. Figure 5 illustrates the SONET multi-layer alarm sequence.

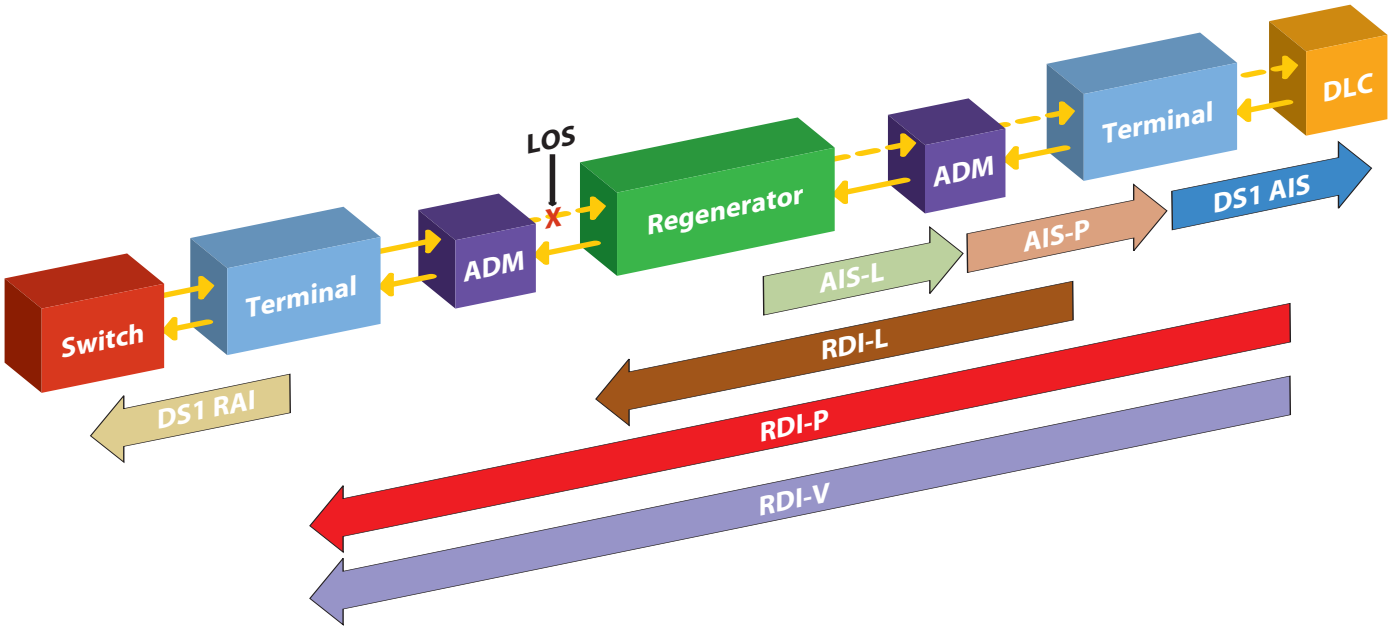

Figure 5: SONET Multi-Layer Alarms

If LOS, LOF or trace mismatch is detected by a section layer NE (e.g., regen), that NE will generate a downstream line AIS. Upon reception of this line AIS, the downstream ADMs and terminals will generate STS path AIS, VT path AIS, and DS-n path AIS on the paths and DS-ns affected by the line failure. Downstream line and path terminating equipment will also generate line and path RDI alarms, RDI-L, RDI-P, RDI-V in the upstream direction. Upstream DSn RAI yellow alarms will also be generated. This complete spectrum of alarms provides appropriate protection switching triggers and enables a downstream NE to squelch the transmission of unnecessary alarms to the OSS.

By virtue of including bit interleaved parity in the overhead, SONET NEs have the ability to detect bit errors at the line, STS path and VT path layers. Detection of such coding violations by a downstream NE, results in the generation of line, STS path and VT path REI in the upstream direction.

At present Ethernet is considerably behind SONET in the area of multi-layer alarms. Currently no standards exist in the area of Ethernet. Work has started on this topic in the ITU-T SG-13 Question 3 Rapporteur Group on Ethernet OAM. ITU SG 13 has also generated recommendation Y.1711 for MPLS OAM [ITU Y.1711], which provides AIS and RDI indications at the MPLS layer. At present. MEF plans to address this topic in the second issue of the MEF Service OAM IA. As will be discussed later, IETF is working on OAM for MPLS tunnels and pseudowires, but to date, this work has not yet included AIS/RDI concepts. Work is taking place on portions of the overall solution in multiple forums, but no single forum has put all the multi-layer components together in a coherent standard like GR-253 for SONET.

Over time, developing multi-layer AIS/RDI for Ethernet services should be possible. The following is a model for how this development could be accomplished. An analogy can be drawn between the Ethernet transport layered model as illustrated in Figures 2 and 3 and the SONET layered model. Figure 6 illustrates how the SONET multi-layer alarm concept can be mapped to E-LAN and E-Line service networks.

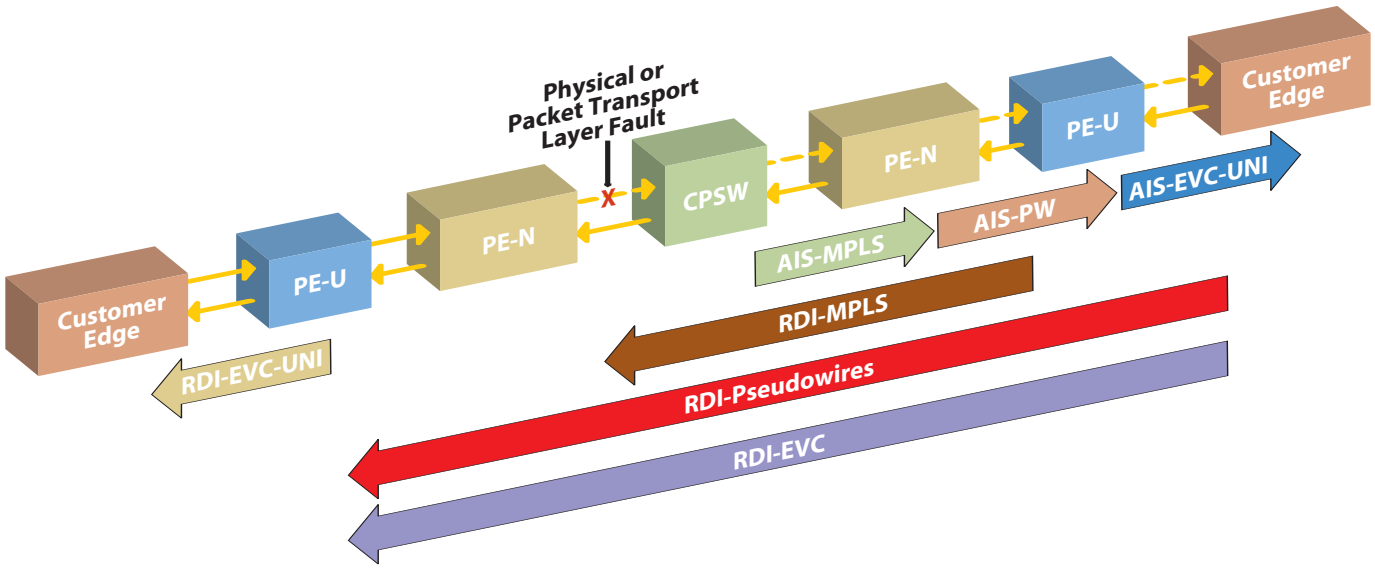

Figure 6: Multi-Layer Alarms for Ethernet

The core packet switch is analogous to a SONET section layer NE (e.g., repeater). If this node detects a fault at the physical or packet transport layer, the node should generate an MPLS tunnel layer AIS (AIS-MPLS) to downstream PE-N for all of the MPLS tunnels impacted by the fault. The PE-N will generate a pseudowire layer AIS (AIS-PW) to the downstream PE-U on all pseudowires impacted by the fault. The PE-U will generate a UNI EVC AIS (AIS-EVC-UNI) on all of the EVCs that are impacted by the fault.

The downstream PE-N will generate an RDI to the upstream PE-N at the MPLS tunnel layer (RDI-MPLS). The downstream PE-U will generate an RDI to the upstream PE-U at the pseudowire layer (RDI-PW) and also generate an RDI to the upstream PE-U at the EVC service layer (RDI-EVC). The upstream PE-U will generate a UNI EVC RDI (RDI-EVC-UNI). These RDIs will be generated for all impacted MPLS tunnels/PW/ EVC.

No direct Ethernet analogy for the SONET REI notifications exist. Errors are detected by using the BIP-8 per frame parity check. Per received frame error, counts are encoded in overhead bytes and are sent upstream as REI indications. At the Ethernet service layer, frame errors are detected by a CRC check. The Ethernet Physical Layer includes a block code, which can check for symbol errors. This process will be accomplished wherever there is a provider bridge function, for example in PE-U and PE-N nodes. No CRC frame check at the MPLS or pseudowire layers exist. Hence, emulating the REI function at the Ethernet service layer, which is equivalent to an REI-V, is the only possibility. This process should be sufficient to support Ethernet services.

Ethernet Link OAM, which is currently being standardized by the IEEE 802.3ah committee, provides error event notifications. These notifications include errored symbol period event, errored frame seconds event, errored frame period event and errored frame seconds summary event. The same error notifications, with the exception of the symbol period event, could be utilized for Ethernet service OAM. When a subscriber-received frame CRC error threshold is crossed, the downstream PE-U could forward one of these notifications to the upstream PE-U.

Performance Management

In order to support Ethernet services, the two principal performance management requirements are SLA verification and network monitoring. SLA verification is accomplished at the Ethernet service OAM layer in PE-U nodes at the egress and ingress to the customer. Per customer EVC performance counters are required. Additional Ethernet service performance statistics can be gathered at the bridge function implemented in PE-N nodes. No CRC checks at the MPLS or pseudowire layers exist. However, per tunnel and per pseudowire performance data can be collected in the form of counts of transmit, receive and dropped packets.

The MEF has recently started to work on an IA for Ethernet service PM [MEF PM]. Their major challenge is to decide how to accumulate sufficient per customer EVC performance metrics in a scalable manner. The most likely approach will be to apply a sampling technique in which a fixed set of performance counters are time multiplexed across the set of customer EVCs, such that maintaining permanent per EVC counters is not necessary. The challenge is to provide adequate SLA verification.

Defining SONET-like PM threshold crossing alerts, which can be a basis for alarms sent to OSS, will also be necessary. The events already defined for Ethernet link OAM can provide an initial basis for these alarms. SONET OAM is ahead of Ethernet OAM in this area. SONET leverages the overhead BIP-8 parity check at section/line/path to provide a rich spectrum of per frame PM statistics and associated alarms.

Ethernet PE - NID Link OAM

Ethernet PE-NID Link OAM monitors and verifies the integrity of the access link between the PE/PE-U and the NID, which is located on the customer premises. The NID is typically a managed media converter. In some cases, the PE-U will be located on the customer premises and will incorporate the NID function. In these cases, this OAM flow is not needed for EE OAM.

Ethernet PE-NE Link OAM utilizes IEEE 802.3ah Ethernet link OAM, [IEEE 802 Link OAM]. This standard, which is close to final agreement, defines OAM capabilities for a single Ethernet link between NEs that terminate the MAC layer and is not forwarded by a bridge. IEEE 802 Link OAM supports all of the key OAM functions required for Ethernet service except traceroute, which is not applicable because the OAM flow is for a single link. Ethernet link OAM frames are distinguished by a slow protocol Ethertype with a sub-type byte indicating OAM and consequently can be sent at a maximum rate of 10 frames per second. The destination MAC address is the link local multicast address that is reserved for slow control protocols. The source MAC address is the port address of the source of the link OAM frames. The information field of the frame includes flag bytes, a code byte which indicates the type of OAM frame and any associated data. The flag field includes bits, which indicate a receive link fault, a critical event (e.g., alarm) and a dying gasp (e.g., unrecoverable local failure). Continuity verification is accomplished by constant exchange of OAM information frames, which include OAM configuration information and a vendor ID. IEEE 802 Link OAM does not support ping but does support

loopback, which is indicated by an OAM loopback frame. With respect to AIS/RDI alarms, IEEE 802 Link OAM supports a receive link failure flag in every frame, which is, in effect, an RDI. AIS is not applicable because the OAM flow is constrained to one link. With respect to REI, IEEE 802 Link OAM generates the following event frames: errored symbol period event, errored frame seconds event, errored frame period event and errored frame seconds summary event.

In addition to the above capabilities, IEEE 802 Link OAM also provides the ability to send/receive vendor-specific OAM frames and to retrieve MIB objects from the downstream NE. Ethernet link OAM is mature in a standards context and is on par with SONET with respect to capability.

Ethernet UNI OAM

The Ethernet UNI OAM flow is between the PE node, which is a UNI DCE, and CE node, which is a UNI DTE. The standards work on UNI OAM is in its very early stages. For SONET, the relevant standards work is associated with the GMPLS control plane. The standards work is more mature than in the Ethernet case as OIF UNI 1.0 provides customers with the status of their circuits. However, if Ethernet and SONET are compared with respect to level of implementation for UNI OAM, both only require minimal implementation.

The MEF plans to address UNI OAM in their UNI IAs. The ITU SG 13 Question 3 Group has also recently started to work on Ethernet OAM, and their current model includes an OAM flow to CE nodes. The current plans for the MEF UNI 2.0 IA include an Ethernet LMI capability. This capability will be similar in concept to the broadly deployed frame relay LMI capability. The E-LMI protocol will provide continuity verification between the PE and CE, as well as a vehicle for the service provider to inform the customer of the status of EVCs and to transmit any relevant alarm information. E-LMI probably will not include loopback capability.

Transport Layer OAM

Transport Layer OAM involves OAM at the MPLS tunnel, pseudowire and IEEE provider bridge Q/Q layers. MPLS tunnel OAM work is currently taking place in the IETF MPLS working group and in the ITU SG 13 Question 3 Group. Pseudowire OAM work has recently started in the IETF PWE3 Group. Initial discussion has been conducted of OAM in the IEEE 802.1 provider bridge group, as well as some early work on L2 VPN OAM in the IETF PPVPV group.

To date, the most mature work is the MPLS tunnel OAM as defined by the IETF [IETF draft-mpls-lsp-ping] and ITU Recommendation Y.1711, which have previously been discussed in this paper. Y.1711 supports connectivity verification and FDI/BDI that is equivalent to AIS/RDI across nested MPLS tunnels. Y.1711 potentially covers pseudowires, which can be implemented by nested MPLS tunnels. Loopback is still under discussion.

IETF draft-mpls-lsp-ping defines how to accomplish ping and traceroute for MPLS tunnels. Unfortunately, the method of distinguishing MPLS OAM packets and the OAM packet format is different between IETF and ITU. ITU uses a reserved MPLS label, whose value was also reserved by IETF to distinguish OAM packets. IETF uses the IETF standard router alert label. IETF uses UDP packets as MPLS OAM probes. ITU uses unique ITU-defined OAM packet format.

An initial draft in the IETF PWE3 Group [IETF draft-nadeau-pwe3-vccv] defines how to implement connectivity verification for a pseudowire. This draft proposed two methods of PW connectivity verification, in band and out of band. The in band approach mandates use of the PW control word, which is sent in addition to the PW MPLS label. The control word includes a flag to indicate an in band OAM channel that exists over the same label as the user plane data. For the out of band approach, an extra router alert label is added to the label stack. For both approaches, the continuity verification ping packet is the UDP packet that is used for MPLS tunnel ping [IETF draft-mpls-lsp-ping]. The draft also includes extensions to the LDP PW signaling protocol to negotiate the OAM transport method between the two ends of the PW.

Ethernet transport layer OAM is functionally equivalent to SONET OAM. At present, the major issue is that the requisite functionality is split across ITU and IETF standards, and differences exist between the two sets of standards. In order to get full functional coverage, both sets of standards will need to be implemented.

Security Management

Security management is an area where Ethernet could have an advantage over SONET in the long term. For SONET networks, security has been supported by implementing a closed DCN network environment and user password control for OSS to NE access. For Ethernet, IEEE 802.1X link authentication already exists, and IEEE 802.1 is close to finalizing IEEE 802.1aa [IEEE Link Security] an updated/improved version.

IEEE Link Authentication enables service providers to verify by strong authentication that an attached customer link is to the correct customer. The value of this capability is two fold: 1) link authentication protects against one customer from stealing service from another 2) link authentication protects against a craft technician inadvertently connecting the wrong customer to a provider port or a valid customer to an incorrect provider port.

IEEE 802 has recently started a Layer 3 security project that will extend Layer 2 Ethernet security standards for additional applications such as EPON and bridge security, including bridging with a secure L2 control plane.

With respect to the MPLS-based Ethernet transport layer, the control and management plane can leverage the full spectrum of already existing IP security standards including IPsec.

Summary and Conclusions

Ethernet has four principal open issues: 1) standards are not complete 2) Multiple standards in different forums are not consistent 3) per EVC performance counters for SLA verification are needed 4) hardware-based traceroute is needed with a smooth migration strategy for existing products to support this capability.

Ethernet offers three principal advantages in comparison to SONET: 1) traceroute 2) better automation 3) enhanced security.

The prognosis (if and when the above issues are addressed) is that Ethernet has the potential to offer enhanced OAM capability in comparison to SONET. Table 1 summarizes the current status of carrier-class Ethernet for the principal OAM flows and functions.

OAM Flow	OAM Function					
	Continuity	Ping/Loopback	Traceroute	Multi-layer Alarms	Performance Management	Security
Service OAM PE-PE	MEF standard in development Automation/ simplicity better than SONET	Ping MEF standard started, delay/jitter also measured SONET = LB Potential Advantage - Ethernet	Standard not started, implementation feasible SONET – not supported Potential Advantage - Ethernet	Standard not started in IETF, early ITU work Positive SONET capability Ethernet can duplicate	MEF standard started Positive SONET capability Ethernet challenge = PM counter scalability	Ethernet has link authentication now + secure management plane + new standard work starting. Advantage Ethernet
PE – NID Link OAM	IEEE 802 Link OAM standard almost compete	LB - IEEE 802 Link OAM standard almost compete	Not Applicable	IEEE 802 Link OAM standard supports multiple link alarms	IEEE 802 Link OAM standard has PM based alarms	Ethernet link security is applicable
UNI OAM	MEF UNI 2.0 ELMI in development OIF UNI 1.0 done - limited SONET implementation	Loopback may not be supported	Not Applicable	ELMI will indicate EVC status OIF UNI provides connection status	Availability of PM based alarms to be determined	Ethernet link security is applicable Can also be supported by OIF UNI
MPLS Tunnel OAM	Supported by ITU standard but not IETF standard	Ping supported by ITU & IETF but different frame format	Supported by IETF standard but not ITU standard	Supported by ITU standard but not IETF standard	MPLS label does not include CRC	Control plane can leverage IP security
Pseudowire OAM	IETF standard started – done by ping	IETF standard started – done by ping	Not Applicable	IETF standard does not address	Pseudowire label or control work does not include CRC	Control plane can leverage IP security

Table 1: Summary Comparison Chart – Ethernet and SONET OAM

References

[1] [MEF Service OAM] Metro Ethernet Forum Approved Draft, "Ethernet Service OAM," Work in Progress, March 2003

[2] [GR253] Telcordia GR-253-Core, "Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria," September 2000

[3] [IETF draft –stokes] Stokes, O et al, "Testing Hierarchical Virtual Private LAN Services," draft-stokes-vkompella-ppvpn-hvpls-oam-01.txt, Internet Draft, December 2002

[4] [IETF draft-mpls-lsp-ping] Kompella, K et al, "Detecting MPLS Data Plane Liveness," draft-ietf-mpls-lsp-ping-02.txt, Internet Draft, March 2003

[5] [ITU Y.1711] ITU Study Group 13, Recommendation Y.1711 "OAM Mechanism for MPLS Networks," July 2002

[6] [MEF PM] Metro Ethernet Forum Approved Draft, "Ethernet Performance Monitoring," Work in Progress, March 2003

[7] [IEEE Link OAM] IEEE 802.3ah Committee, "Draft Amendment to IEEE Std 802.3-2002 Ethernet in the First Mile Draft D1.414," Work in Progress, April 2003

[8] [IETF draft-nadeau-pwe3-vccv] Nadeau, T. et al, "Pseudo Wire (PW) Virtual Circuit Connection Verification (VCCV)" draft-nadeau-pwe3-vccv-00.txt, Internet Draft, February 2003

[9] [IEEE Link Security] IEEE 802.1 Committee, "Draft Standard – Port Based Access Control Amendment 1: Technical and Editorial Corrections Draft D5," Work in Progress, February 2003

Acronym	Descriptor
ADM	Add Drop Multiplexer
AIS	Alarm Indication Signal
BDI	Backward Defect Indicator
BIP	Bit Interleaved Parity
CE	Customer Edge
CPSW	Control Point/Switch
CRC	Cyclic Redundancy Check
DCC	Data Communications Channel
DCE	Data Communications Equipment
E-LAN	Ethernet Private LAN
E-Line	Ethernet Private Line
E-LMI	Ethernet Link Management Interface
EMS	Element Management System
EoS	Ethernet over SONET
EPON	Ethernet Passive Optical Network
EVC	Ethernet Virtual Connection
FDI	Forward Defect Indicator
GFP	Generic Framing Procedure
GMPLS	Generalized Multiprotocol Label Switching
IA	Implementation Agreement
ICMP	Internet Control Message Protocol
IEEE	Institute of Electrical and Electronics Engineers, Inc.
IETF	Internet Engineering Task Force
IP	Internet Protocol
ITU	International Telecommunications Union

Acronym	Descriptor
L2TP	Layer 2 Tunneling Protocol
L2 VPN	Layer 2 Virtual Private Network
LAN	Local Area Network
LDP	Label Distribution Protocol
LMI	Link Management Interface
LOF	Loss of Frame
LOS	Loss of Signal
MAC	Medium Access Control
MAN	Metropolitan Area Network
MEF	Metro Ethernet Forum
MIB	Management Information Base
MP EVC	Multipoint Ethernet Virtual Circuit
MP2MP	Multipoint-to-Multipoint
MPLS	MultiProtocol Label Switching
NE	Network Element
NID	Network Interface Device
NMS	Network Management System
O/E	Optical/Electrical
OAM	Operations, Administration and Maintenance
OIF	Optical Internetworking Forum
OSS	Operations Support System
P2P	Point-to-Point
PDH	Plesiochronous Digital Hierarchy
PE	Provider Edge
PE-N	Provider Edge–Network

Acronym	Descriptor
PE-PE	Provider Edge to Provider Edge
PE-U	Provider Edge–User
PM	Performance Monitoring
PP	Point-to-Point
PP EVC	Point-to-Point Ethernet Virtual Circuit
PPP	Point-to-Point Protocol
PPVPN	Point-to-Point Virtual Private Network
PW	Pseudowire
Q/Q	IEEE 802.1Q Tunneling
RAI	Remote Alarm Indication
RDI	Remote Defect Indication
REI	Remote Error Indication
RMON	Remote MONitoring
RSVP-TE	Resource ReSerVation Protocol–Traffic Engineering
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
STS	Synchronous Transport Signal
TDM	Time Division Multiplexing
TMN	Telecommunications Management Network
TTL	Time To Live
UDP	User Datagram Protocol
UNI	User Network Interface
VLAN	Virtual Local Area Network
VPN	Virtual Private Network
VT	Virtual Tributary