White Paper

Ethernet OAM

By Yaakov (Jonathan) Stein, Chief Scientist
RAD Data Communications

RAD
data communications
Innovative Access Solutions

Abstract

Until recently Ethernet lacked OAM functionality like that found in in SONET/SDH or ATM, and therefore was not characterized as “carrier class.” Now it has acquired two types of OAM, one developed by the ITU and IEEE 802.1, and another by the EFM task force. The former is a full-featured OAM mechanism that can run end-to-end and includes all traditional OAM elements such as AIS, RDI and performance measurement. The latter is limited to continuity monitoring of a single link and is targeted at access applications. As the former protects the Ethernet service layer and the latter the Ethernet physical layer, the two can be complementary. However, it is not clear why there is a need for two different standards.

Contents

Introduction

OAM (Operation, Administration, and Maintenance) describes the monitoring of network operation by network operators. OAM is a set of functions used by the user that enables detection of network faults and measurement of network performance, as well as distribution of fault-related information. OAM may trigger control plane or management plane mechanisms, e.g. by activating rerouting or by raising alarms, but such functions are not part of the OAM itself. OAM functionality ensures that network operators comply with QoS guarantees, detect anomalies before they escalate, and isolate and bypass network defects. As a result, the operators can offer binding service-level agreements. The operation of networks without OAM demands many more resources for continuous manual intervention to detect failures, expensive truck rolls to localize faults, and human performance measurement. These networks have lower availability and longer down times, and are more expensive to maintain.

OAM is generally utilized for detecting and localizing network faults, examining and reporting network status, monitoring network performance, and provisioning and configuring user parameters. Traditional TDM systems provide OAM fields in their overhead, while modern packet-switched network OAM protocols *may* exploit packet overhead to piggyback OAM information onto user traffic, but more frequently introduce special-purpose OAM packets alongside user packets.

Each layer of a layered network needs its own OAM. For example, a VoIP application run over an IP network and layered on an ATM network supported by a SONET infrastructure needs four types of OAMs:

- Application layer OAM for the VoIP based on RTP mechanisms
- IP layer OAM (usually limited to troubleshooting by ping and traceroute)
- ATM layer OAM
- SONET OAM (at the bottom of the stack).

By keeping these OAM mechanisms distinct, end-users, service providers and network operators can each monitor the status of functions for which they are responsible, and make corrections without involving others. While defects on one network layer will usually impact higher layers (such as server layer defects which cause loss of client layer information), faults detected on one layer do not necessarily imply correlated faults on lower layers.

OAM mechanisms were first devised for PDH networks, and included AIS and RDI components (see below). PDH OAMs became more sophisticated, and with the advent of SONET/SDH OAM became a crucial network component. TDM OAMs supported the famous “five nines” availability and reliability, and 50 millisecond recovery times from failures. ATM included fault detection and performance measurement OAMs from its initial design stages, and MPLS has several flavors of OAM developed by the IETF and by the ITU. Ethernet is the most recent communications protocol to adopt OAM capabilities.

OAM for Ethernet

Ethernet originated as a Local Area Network (LAN) technology. Since LANs usually consist of a relatively small number of co-located stations, all managed by a single entity, defect detection was manual, and performance was never really a concern. As Ethernet developed, physically separated LANs were interconnected but still managed by a single entity (although now an enterprise). OAM was still not a major concern, and network defects were handled by manual activation of simplistic tools such as ping.

Since the introduction of Metro Ethernet Networks and the advent of “Carrier Class Ethernet” the situation has changed radically. These networks need to be managed by service providers, and in order to be truly “carrier class” it is essential for Ethernet MANs to support automated defect detection and performance measurement. In order to guarantee SLAs, service layer parameter monitoring is also required. IP-based tools such as ping and traceroute are not suitable for pure Ethernet networks, and even when IP tools are being run they function at a higher layer, and thus do not directly relate to the underlying Ethernet network or service.

In order to enable Ethernet service providers to operate and maintain their networks, there is a need to include OAM on the Ethernet layer. This new OAM must integrate seamlessly with existing Ethernet protocols in order to encourage its adoption while enabling coexistence with conventional non-OAM-capable Ethernet devices.

Belatedly, two Ethernet OAM protocols have emerged. One has been developed for “Ethernet in the first mile” (EFM) applications, operating at the level of the single link, while the other tackles the wider problem of end-to-end Ethernet connectivity and service guarantees. The link-layer OAM was developed by the 802.3ah EFM task force in the IEEE 802.3 working group, and thus is often called the “802.3ah” or “EFM” OAM. This task force's work was incorporated into the main 802.3 Ethernet specifications as Clause 57. We will refer to this OAM as the "Ethernet link-layer" OAM; service-layer OAM has not been completely specified. In the IEEE the work is being conducted under the name 802.1ag (Connectivity Fault Management), while in the ITU-T the draft Recommendation was known as Y.17ethoam, although it is now called Y.1731. The Management Area of the MEF is also working on "Ethernet Service OAM.." Fortunately, these three organizations are cooperating with substantial participant overlap, and it is expected that the final output will be sanctioned by all three SDOs. We shall refer to this new OAM as "Ethernet service OAM."

Due to their different objectives, Ethernet link-layer OAM and Ethernet service OAM protocols were not intended to compete with each other, and can even be complementary. In an access segment based on EFM, both Ethernet OAM protocols may be running simultaneously, where the EFM OAM monitors the lower (physical transport) layer, and the service OAM maintains the higher layer. However, implementing two very different standards will be cumbersome to both equipment vendors and operators. As service OAM can be restricted to a single link and contains a superset of the link layer functionality, deployment of EFM OAM will probably be limited to simple first mile applications.

OAM Functions

OAM was originally developed by network operators to increase the reliability and maintenance of TDM networks. The OAM mechanisms were so successful in reducing network maintenance costs that it was natural to extend them to modern frame and packet-based networks.

OAM is a user-plane protocol, although it may influence the operation of control-plane and management-plane functions. For example, OAM is not directly responsible for protection switching, but may be trigger this by detecting a defect in the network that necessitates switching. In addition, OAM does not provision links, but it can signal provisioned links are up and ready for use. In order for OAM mechanisms to accurately share the fate of user traffic, they need to follow the same path as user data.

The primary function of an OAM protocol is to detect network defects. The formal term for the smallest detectable discrepancy between observed and expected operation is *anomaly*. Isolated anomalies, such as correctable bit errors or timing deviations within certain limits, do not interfere with network operation and thus do not need to be reported. However, when there is a sufficient density of anomalies (i.e. a large enough number of anomalies in a given time period) some desired function may be impaired, resulting in a *defect*. Usually a *single fault* cause yields many different correlated defects. When the fault cause persists for long enough there is a *failure*, meaning that a network function is terminated. An *alarm* is an indication to humans that there was a failure. Alarms do not sound for momentary glitches, but only for defects that persist for long enough (a matter of seconds).

Since TDM networks transfer data at a constant rate, any discontinuity is immediately identified as a Loss of Signal (LOS) defect. For frame-based or packet-based networks, discontinuities may not be as immediately apparent, as it is not known when a frame or packet is expected. To remedy this situation an OAM message known as Continuity Check (CC) may be employed. This message is also known as Connectivity Check (with the same acronym) or Connectivity Verification (CV), although these terms are more properly reserved for verification of proper interconnection topology. CC is implemented by periodically sending messages across the network to detect continuity failures; if the remote network element does not receive the CC message in a timely fashion, it (the remote network element) discerns a loss of continuity (LOC) defect. Note that CC is unidirectional, as distinct from the "ping" packets used in IP networks that must be returned from the remote network element for the originator to detect a defect in the bidirectional path.

CC messages are an extremely efficient way of monitoring network connectivity, and thus there is little reason for them not to be universally employed. They may be multicast from a single source, thus obviating the N^2 message streams to verify connectivity of a network with N network elements by pinging. Furthermore, CC messages may be sent at a relatively slow rate (e.g. once every second), thus minimizing their impact on network bandwidth utilization.

A TDM switch detecting the persistence of LOS or other serious defects performs several actions. First, it raises an alarm in order to notify the operator's personnel of the problem. Second, it replaces the missing data with a standard bit pattern (usually all 1s) known as AIS and transmits this artificial data to the next switch along the path. AIS originally stood for Alarm Inhibition Signal, and its function was to inhibit the raising of alarms at all successive switches, so that the network operator would only receive an alarm indication from the affected switch. However, AIS came to be renamed as Alarm Indication Signal, although this term is not entirely accurate; the alarm is not indicated, rather the defect or failure is indicated and the alarm is sounded.

AIS is an OAM message that advertises that a defect has been detected at some previous point along the path. Forward Defect Indication (FDI) is a more general message that indicates to following network elements that a defect has been detected at some previous network element along the path. For bidirectional network connections, Reverse Defect Indication (RDI) reports on a unidirectional defect in the reverse direction, such as a destination node could not receive traffic. RDI is generally employed when only portions of the network are directly managed. Unmanaged remote sites cannot directly raise an alarm, so unidirectional loss of connectivity to a managed site is indicated using RDI.

The above OAM messages operate, either by design or due to defects, during normal operation of the network. However, there is another class of OAM messages that is used during the initial setup of network equipment, or after defects have been detected, to hone in on possible or putative problems. The most important of these are loopback (LB) messages. For TDM systems, loopback always meant that an interface is placed into a diagnostic mode whereby its input was immediately sent to its output. This could be done manually, or by sending an OAM loopback *command* to the device. In this way fault causes could be localized from the source by successively placing each network element along the path into loopback mode.

Loopback messages for frame and packet-based networks may be divided into two general categories:

- I*n-service* (also called *intrusive*) and
- *Out-of-service* (*non-intrusive*).

In-service LBs are OAM packets that can be sent without disrupting the normal operation of the network. When a device receives an in-service LB OAM packet with its address, it is returned to the sender rather than forwarded; all other packets are forwarded as usual. Ping messages are in-service loopback messages of this sort, and form the basis of diagnostic and performance measurement toolkits of IP and MPLS networks.

On the other hand, an out-of-service LB command is an instruction to the interface to go into a special loopback mode that disrupts its normal operation. All packets (except OAM packets themselves, since if they were disrupted there would be no way to end the LB mode) are immediately returned to their source without further handling. If the source does not receive the loopback "echo" within a certain time, it realizes that there is a defect in at least one of the directions between itself and the device placed in LB mode. Note unlike the simple "short-circuiting" performed by TDM systems, packet-based devices must be processed in order to implement loopback functionality. At the very least, the source and destination addresses of the received packet must be reversed.

Some OAM protocols support both in-service and out-of-service loopback, since the two fulfill different functions. Out-of-service loopback is frequently used while setting up a connection before the service is provided, and may be used to measure throughput of unloaded paths. It may also be used as a last option in troubleshooting a failed service. On the other hand, in-service loopback is a simple method of continuously monitoring bidirectional continuity in order to rapidly detect service failures and trigger protection switching mechanisms.

Related to in-service LB, but more informative, are *trace-route* (or *link-trace*) mechanisms. Not only do these check bidirectional continuity between two points, they also identify the network forwarding elements that lie along the path between them. The IP traceroute mechanism is based on ping packets with hop-based timeouts. The traceroute originator sends a ping set to timeout at the next hop, and the return message identifies that adjacent router. Next, a ping packet is sent set to timeout after two hops, and so on.

While both TDM and packet-switched networks can suffer from faults, there are inherent differences between the two network types. TDM networks are inherently predictable and stable. Absent a failure, every bit sent from the source is received by the destination (although in rare instances it may be received incorrectly), and delay along a path is constant throughout the life of the path. This is not the case for packet-switched networks, where packets are routinely lost due to errors, congestion, or policy decisions. Packets or frames are received by the destination after highly variable packet delay variation (PDV) or frame delay variation (FDV). Hence, for frame and packet-switched networks, OAM mechanisms for measurement of network performance are important. Performance parameters typically measured include loss, one-way delay, two-way delay, and delay variation.

How can such parameters be measured? To do so, only minor extensions are needed for the mechanisms previously described. For example, CC-type messages can be used for approximate packet-loss determination by adding a sequence number and functionality for keeping track of the percentage lost. By further adding timestamps, CC messages can be used to measure one-way delay, and non-intrusive LB messages can be used to measure two-way delay. By keeping track of delay measurements, delay variation can be calculated. CC and LB messages can be similarly used for *throughput* determination in idle network segments, and even for networks in use by using more sophisticated statistical sampling techniques.

A final capability sometimes appended to OAM systems, although perhaps not technically an OAM function, is configuration management. Provisioning end points from a central location is obviously a management function, but retrieval and configuring parameters of remote terminals is sometimes accomplished via OAM messaging, especially in systems that lack fully developed control protocols.

EFM (802.3ah) Link-layer OAM

As mentioned before, Ethernet link-layer OAM was developed for “Ethernet in the First Mile” (EFM) applications. The expression *last mile* is frequently used for access links reaching customer locations from an operator’s central office. If that is the case, why did the IEEE decide to reverse this terminology and call Ethernet access *first mile*? The reasoning is easy to understand. From the perspective of *core network* engineers, that last portion of the network reaching the customer is obviously the “last mile.” On the other hand, *Ethernet* originated as a customer-controlled local network, so it is natural to think of the first portion of its extension to a wider area network as the “first mile.”

Core networks have rich connectivity, with multiple alternative paths available between any two switches or routers. In contrast, first mile networks usually have very simple topology, direct point-to-point links between provider and customer, or point-to-multipoint PONs so that the provider can broadcast content to many customers. For this reason, the OAM needed for EFM applications is link-layer OAM.

The capabilities of link-layer OAM are limited, being restricted to placing the remote device into loopback, setting flags indicating critical events, and querying the remote device’s configuration. There is no performance measurement and the information exchanged about the state of the link being monitored is minimal. More significantly, since this OAM is limited to a single link there can be no AIS indication of failure of a previous path segment and thus no end-to-end service guarantee.

IEEE link-layer OAM operates purely at the Ethernet layer, and so (unlike SNMP or ping) does not require an IP address. This means that Ethernet service providers don’t need to run IP protocols or manage IP addresses. Furthermore, special Ethernet features may be directly supported, such as Ethernet multicast and slow protocol frames. When an OAM frame is received by an OAM-enabled Ethernet MAC, it is passed to the OAM client for processing; such a frame is simply discarded if received by a MAC that does not support link-layer OAM. In any case, link-layer OAM frames are never forwarded.

Since the IEEE link-layer OAM is generally used over a link between a service provider and a customer, it defines two modes for OAM entities: *active* or *passive*. The elements of the provider network (e.g. DSLAMs or provider Ethernet switches) operate in active mode, and can exert control over the passive-mode devices (e.g. DSL modems or customer premises switches). Thus, the active-mode entity can send an LB command forcing the passive-mode device into loopback mode, and query the configuration parameters of the passive-mode device. However, the reverse is not possible. We shall see later that service OAM specifies more complex relationships, involving end points, intermediate points, and levels of hierarchy.

Link-layer OAM messages are sent in untagged *slow protocol* frames called OAM Protocol Data Units, or OAMPDUs. Slow protocols are protocols used to control operational characteristics of the Ethernet device, such as the Link Aggregation Control Protocol (LACP, formerly known as 802.3ad) which also utilizes slow protocol frames. Slow protocols are slow in the sense that they are restricted in the number of protocol frames that may be transmitted per second (for OAMPDUs – no more than 10 frames per second), thus facilitating software implementations of the OAM client. All slow protocols use Ethertype 88-09, and link-layer OAM is differentiated by a sub-type of 03 that appears as the first byte of the MAC client payload. OAMPDUs are multicast to a specific multicast address that is *link-constrained*, since OAMPDUs only traverse a single link, but are never forwarded by bridges or switches, even if these bridges do not implement OAM.

OAMPDUs contain control and status information needed to monitor, test and troubleshoot OAM-enabled links. This information is encoded using a major code followed by information encoded in Type-Length-Value (TLV) format (see Figure 1). Many modern protocols use TLVs in order to permit protocol extensions. Proprietary (organization specific) OAM extensions may support enhanced capabilities, but will still be limited by the single-hop nature of EFM OAM.

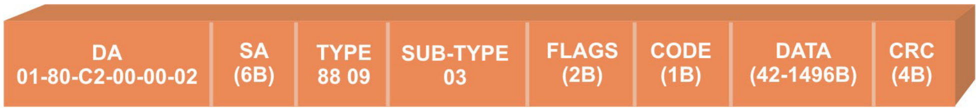

Figure 1 - EFM OAM Format

The best way to explain the capabilities of link-layer OAM is to enumerate the flags and codes. Six codes are presently defined:

- Information
- Event notification
- Variable request and response
- Loopback control
- Organization-specific.

Information code OAMPDUs are used for autodiscovery, heartbeat, and fault notification, making them the workhorse of link-layer OAM. Discovery is the procedure whereby OAM-enabled entities discover each other and exchange information regarding their OAM capabilities and configuration. The OAM capabilities may be used to determine whether it is worthwhile to run the OAM protocol. For example, this may be used for active-mode OAMs interested only in loopback-capable peers. OAMPDUs must be sent at least once per second; if there are no other pending messages, an information PDU with no information TLVs is sent. Finally, although all OAMPDUs of all codes carry flags (see below), when an event occurs that necessitates the immediate sending of a fault indication, an information code OAMPDU is always sent.

Event notification OAM frames are used for reporting various link statistics such as the number of symbol errors that occurred during a specified period, or the running total of frames with errors since the OAM sublayer was last reset. *Variable request* frames are used by service providers to obtain the customer's configuration by requesting MIB variables. *Variable response* frames are returned by the customer in response to such requests. Only active-mode entities may send request frames, and only passive-mode ones will return response frames. Since delivery of any given Ethernet frame is not guaranteed (especially during problematic network conditions when OAM is most valuable), event OAMPDUs may be sent several times to increase the likelihood of reception.

Loopback control OAM frames are used by an active-mode OAM entity to enable or disable intrusive loopback in the remote passive-mode device. The loopback function is an optional feature that may be implemented in software or hardware. When enabled, all frames except OAMPDUs and PAUSEs are returned to their source on the same port instead of being forwarded as usual. At any time during loopback mode, statistics from local and remote OAM clients may be queried.

Flags are contained in every OAMPDU in order to expedite notification of critical events and to continuously monitor for slowly deteriorating performance. These flags signal whether the remote OAM entity has completed the discovery phase and is ready to receive OAM messages. The defects that can be indicated are link fault, dying gasp, and other unspecified critical events. The link fault flag is set when the PHY receiver detects a loss of signal, making this flag a reverse defect indication (RDI). For the duration of the link fault, this flag remains set in all OAMPDUs, and empty information OAMPDUs are sent once per second. Dying gasp is a notification sent by a device about to operationally fail (due to power failure, having been reset, etc.). It originated in the xDSL modem world, and can be used to inform the service provider of imminent failure of the remote device. Any other error condition that does not result in rebooting of the device is indicated by the critical error catch-all flag.

Y.1731/802.1ag Service OAM

Link-layer OAM allows detection of faults on an EFM link. While this is a useful feature, service providers are in need of the ability to fully monitor a customer's end-to-end Ethernet service. This monitoring ability should be agnostic to the layers supporting this service, which may be EFM (i.e. DSL or native Ethernet fiber), but may also be other services such as SONET, ATM, or MPLS ("Ethernet pseudowires"). In addition, when several service providers and/or network operators are involved, each needs to separately monitor the layer for which it is responsible.

When an Ethernet service is first initiated, end-to-end path integrity needs to be verified. Similarly, if some failure is detected, the service provider needs to quickly identify the customers that are affected, and what rerouting can be performed. These functions could be done using IP-based tools, but (as discussed above) if the service being provided is Ethernet, the best approach is to have independent Ethernet-specific OAM.

Ethernet service OAM defines a hierarchy of up to eight OAM levels, allowing users, service providers and operators to run independent OAMs at their own level. By default, users are allocated three levels, service providers two levels, and operators three levels. OAM frames belonging to higher levels are transparently forwarded by lower level devices (e.g. user OAM frames are forwarded by service provider switches).

Service OAM for Ethernet builds on general principles developed over time for TDM, SONET, ATM, and most recently MPLS networks, but Ethernet presents several new challenges. First, previous packet-based OAM protocols were for connection-oriented networks, while Ethernet is connectionless. Hence mechanisms that protect well-defined paths need to be extended to connectionless flows. Second, previous OAM protocols were mainly developed for point-to-point connectivity, while Ethernet is inherently multipoint-to-multipoint. Third, even among connectionless networks Ethernet's behavior can be mischievous, at times sending frames to unexpected places. This is problematic since OAMs should not leak from one OAM domain to another, and customers (or other service providers) should not be able to interfere with proper service OAM operation.

Solving all of these problems requires defining a number of new concepts, and although the IEEE and ITU don't always give these concepts the same names, they are in agreement as to their intentions. Since Y.1731 has been finalized while 802.1ag is still in progress, we shall mainly use ITU terminology here. Y.1731 defines a maintenance entity (ME) that requires management. Some examples of these entities are the entire Ethernet network between two customer switches, or the Ethernet network in the administrative domain of a single service provider, or even a single Ethernet link. Thus MEs can be nested, with link MEs internal to service provider, and MEs of successive providers internal to the customer end-to-end ME.

In order to capture the multipoint-to-multipoint nature of Ethernet, MEs are grouped into ME groups (MEGs, referred to as Maintenance Associations or MAs in IEEE language). A multipoint-to-multipoint Ethernet network with N end-points has N(N-1)/2 MEs, while a point-to-point connection has only one. In order to enable detection of incorrect connectivity, each MEG is given a unique ID, and OAM messages specify the MEG ID for which the message is intended.

At the ends of managed entities we find MEG End Points (MEPs), which are the functions that generate and process OAM frames to monitor and maintain the ME. There may also be MEG Intermediate Points (MIPs) that can respond to OAM messages, but cannot originate them. For point-to-point MEGs, a MEP has a single peer MEP, but in between there may be many MIPs. Hence a MEP can send CC messages to its peer MEP, or direct non-intrusive LB messages towards the peer MEP or to any MIP. It is the responsibility of the MEP to prevent OAM messages from leaking out of the administrative domain to which they belong, or entering another domain. However, MEPs transparently pass OAM frames from other domains when they belong to a higher OAM level, thus enabling end-to-end management of customer connectivity (see Figure 2).

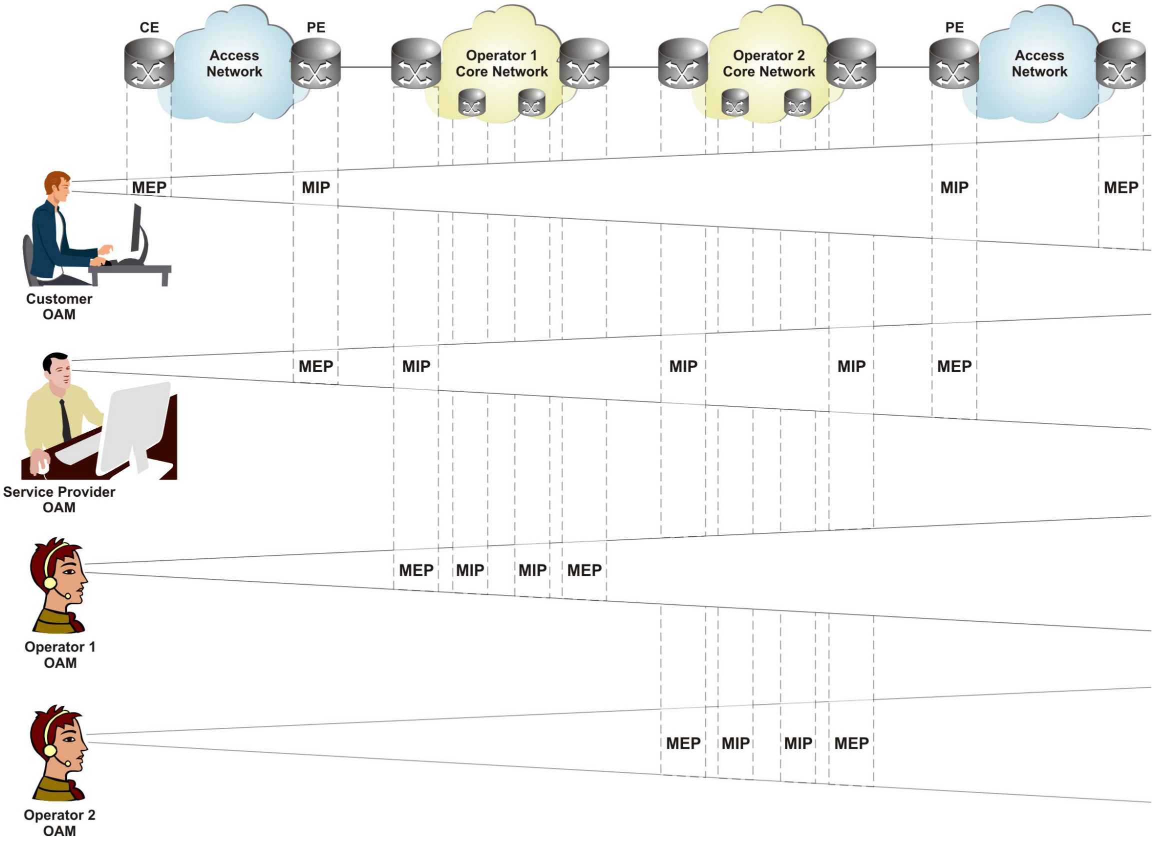

Figure 2 – Service OAM MEPs and MIPs

Y.1731 supports an impressive array of OAM messages, including CC, LB, Link Trace (LT), AIS, RDI, Lock Signal (LCK), Test Signal (Test), Automatic Protection Switching (APS), Maintenance Communications Channel (MCC), Experimental (EXP), and Vendor Specific (VSP) for fault management, and Loss Measurement (LM) and Delay Measurement (DM) for performance monitoring. CC is defined as a *proactive* OAM message, which means that once started it is automatically generated at a configured rate, while other messages are *on-demand*. Yet other OAM message types, such as AIS, are event-triggered but may be enabled or disabled. We will discuss a few of the most important of these messages after detailing the OAM frame format.

After using the DA, SA and EtherType (which will be allocated by the IEEE towards the end of the 802.1ag work), and optionally VLAN tagging, all Y.1731 PAM PDUs use the header depicted in Figure 3.

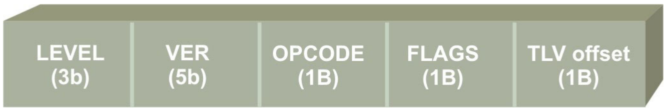

Figure 3 - Service OAM header

The level specifies to which of the eight OAM levels the OAM frame belongs by default:

- 0 through 2 are for customers
- 3 and 4 are for providers domain, and
- 5 through 7 identify an operator domain.

Note that level 0 is the *highest* level and level 7 is the *lowest*. The version number is set to zero. The operating code identifies the message type, and the flags add additional information, such as RDI indication and the CC period. The TLV offset tells us how many bytes must be skipped over before finding the TLV information (for example, if the offset is zero, the TLVs commence immediately after the header). TLVs may not start immediately after the header of Figure 2 because some OAM types require non-TLV encoded information, such as sequence numbers or timestamps. These types of information are commonly set and read by hardware, and thus are best placed at known offsets from the beginning of the OAM PDU, rather than contained in a TLV.

Following the header depicted in Figure 2, the OAM PDU contains a series of TLVs, terminated by a special “end TLV.” For example, the MEG ID is encoded as a TLV, as are the data and test patterns used in LB messages. All Y.1731 TLVs have one-byte type fields and two-byte length fields. There may, or may not, be a value field. The end-TLV has type set to zero, and has no length or value fields.

Arguably the most important of the OAM messages defined by Y.1731 is the CC message. Once enabled, a MEP sends a CC “heart-beat” messages periodically at one of seven possible transmission periods:

- 300 per second
- 100 per second
- 10 per second

- 1 per second
- 1 per 10 seconds
- 1 per minute, or
- 1 per ten minutes.

MIPs transfer CC messages transparently. The CC message enables detection of loss of continuity between MEPs, and an MEP declares a loss of continuity when it does not receive the expected CC messages for 3.5 times the configured transmission period (as indicated in the flags field). However, CC messages also perform various other defect and performance monitoring activities, such as the discovery of other MEPs (by sending a CC to a multicast address), detection of unwanted connectivity between MEPs, and RDI (indicated by setting the MSB of the flags field).

Loopback Messages are transmitted by an MEP upon demand by the administrator to verify bidirectional connectivity to a particular MIP or MEP. Similar to ping messages, LBs are sent upon demand and thus may be sent once, repetitively, or according to any other scheme dictated by the initiator. LB messages can be unicast to the MAC address of the desired entity, or multicast to all peer MEPs in the MEG. The unicast version infers connectivity by timely receipt of an LB response message, while the multicast version produces a list of MEPs with which connectivity was detected. LB messages may optionally carry test patterns, the length and content of which may be configured according to need.

While the LB functionality does not identify the path between OAM entities, the link trace (LT) message does. LT produces the sequence of MIPs and MEPs crossed when forwarding from the source MEP and a given MEP or MIP, and can thus also be used to detect link failures or loops.

AIS messages are used to suppress alarms at client layers from defects discovered at server layers. These messages are not required when the Ethernet network employs the spanning tree protocol (STP), as STP has its own mechanisms for this purpose. When an MEP configured for AIS detects a fault condition, e.g. by not receiving CC messages, it starts transmitting AIS messages periodically at a client MEG Level. A client level MEP receiving AIS suppresses alarms associated with all peer MEPs.

Another message type is the lock message, or LCK. It is used to inform MEPs of intentional diagnostic actions, enabling client MEPs to differentiate between the defect conditions and intentional actions at the server layer that may result in traffic disruption. The test signal, or TST, is used to perform one-way in-service or out-of-service diagnostics tests, verifying throughput, frame loss, errors, etc. The automatic protection switching message APS is used to control protection switching operations, and the maintenance communication channel MCC provides a maintenance communication channel between a pair of MEPs in order to make remote maintenance possible. Experimental EXP OAM messages may be used for trying out new OAM ideas within an administrative domain, and vendor-specific OAM VSPs may be used by vendors between their own equipment.

Summary

While service layer OAM provides the fault detection and performance monitoring for Ethernet services, lower network layers are often still required to run their own OAM. For example, Y.1731's CC mechanism may detect a loss of continuity on an end-to-end service supported by an Ethernet pseudowire over an MPLS network or an ATM network, and invoke the lower layer's native OAM to rapidly identify the fault's cause.

However, when both the service and transport infrastructure is based on Ethernet, the question arises as to whether two independent and incompatible OAM systems are needed. In purely theoretical terms, the Ethernet service layer (called the ETH layer in ITU documents) is a distinct entity that just happens, to share the same format as the transport layer (called the ETY layer). However, this coincidence should not impact on OAM deployment. It is clear that service OAM can be used in a degenerate form over a single link, and the functionality provided is a superset of that obtained by using EFM's OAM (with the exception of the EFM protocol's capability of retrieving remote configuration parameters). Furthermore, it would seem that link-layer OAM, although said to be suitable for all Ethernet links, is going to be deployed almost exclusively in EFM scenarios, where the connectivity is simple and the switches have only very limited computational resources.

From the points of view of both the customer and the Ethernet service provider, the service layer OAM is the most important OAM protocol. Service layer OAM provides clear demarcation points between customers and service providers, between providers, and between providers and network operators. Y.1731/802.1ag is a fully-functional state-of-the-art OAM protocol, and its design utilizes the experience of IEEE and ITU in both Ethernet and carrier-class OAM. Although present implementations are still pre-standard, once can expect fully compliant and interoperable implementations to become available in the near future.

www.rad.com

Corporate Headquarters
RAD Data Communications Ltd.
24 Raoul Wallenberg Street
Tel Aviv 69719, Israel
Tel: 972-3-6458181
Fax: 972-3-6498250
email: market@rad.com
www.rad.com

U.S. Headquarters
RAD Data Communications Inc.
900 Corporate Drive
Mahwah, NJ 07430 USA
Tel: (201) 529-1100,
Toll Free: 1-800-444-7234
Fax: (201) 529-5777
email: market@radusa.com
www.radusa.com

RAD
data communications
Innovative Access Solutions

Catalog no. 802354 Version 3/2006