Q & A

# Enabling the “Experience Provider” Transition at the Network Layer

This Q&A discusses several technologies available in the network layer that boost the quality and efficiency of multimedia service delivery, including Call Admission Control for IPTV and video on demand (VoD), service assurance (including Ethernet OAM), and the Wireless Services Module (WiSM) on the Cisco® 7600 Series routers.

## CALL ADMISSION CONTROL FOR IPTV AND VoD

**Q. What is Call Admission Control (CAC)?**

**A.** CAC prevents oversubscription of video and voice-over-IP (VoIP) networks. It typically applies only to video and voice traffic and not to data traffic. CAC mechanisms extend the capabilities of quality of service (QoS) tools to protect video traffic, in particular from the negative effects of other video traffic, and to keep excess video traffic off of the network. CAC is used for congestion control within video networks. It is a preventive congestion-control procedure and is used in the call-setup phase. CAC can be used to prevent congestion in connection-oriented protocols such as ATM as well as in IP networks.

**Q. What is the Cisco CAC solution for VoD and Broadcast TV?**

**A.** Given the bandwidth demands and complexity of video services, Cisco Systems® has focused significant efforts on developing an Integrated Video Admission Control solution. Integrated Video Admission Control consists of two highly reliable components: Cisco VoD Connection Admission Control and Cisco Broadcast Video Admission Control.

- **Cisco VoD Connection Admission Control** includes an “on-path” (on the data path) signaling component that performs admission control by communicating between the VoD server and the Cisco routers to determine the bandwidth available in the complex and dynamic metropolitan aggregation network topologies. A second “off-path” component uses a policy server to examine the capacity of the access link to a subscriber’s home or the subscriber’s account policies and to communicate back to the VoD server whether or not to admit the stream.
- **Cisco Broadcast Video Admission Control** allows service providers to control oversubscription of broadcast video, reducing the amount of bandwidth consumed by a broadcast without degrading the user experience through use of a connection-admission algorithm. Used in conjunction with IP multicast, which provides intelligent forwarding of IP video streams to ensure the best use of available bandwidth, the admission-control algorithm runs in the aggregation router.

Admission control is a crucial technology for service providers deploying VoD and broadcast video, both of which can impact the total amount of bandwidth available and, if uncontrolled, can lead to a poor-quality user experience. Cisco has been developing the ICAC solution for three years and leads the industry in this innovative and flexible solution. Integrated Video Admission Control can differentiate between paid VoD, free VoD, specific high-demand broadcast channels, and many other variables, giving providers and customers enhanced choices and control over the network video experience.

**Q. What are the typical QoS requirements for voice and video?**

**A.** Table 1 summarizes these typical QoS requirements.

**Table 1.** QoS Requirements for Voice and Video

| Service | Bandwidth Required | Allowed Drop Rate | Allowed Jitter |
|---|---|---|---|
| High-speed Internet | 5 Mbps | | |
| VoIP | <0.2 Mbps | 10** –2 | ~ = 60 ms |
| MPEG 2 video | 3.75 Mbps | 10** –6 | ~ = 200 ms |
| MPEG 4 video | 2 Mbps | 10** –6 | ~ = 200 ms |
| High-definition TV | 6–15 Mbps | 10** –6 | ~ = 200 ms |

**Q. Why is CAC needed for VoD or IPTV?**

**A.** Operators typically architect their networks based on peak concurrency. For example, on a worst-case, Saturday-night scenario, 20 percent of all homes will be viewing one standard-definition stream. But certain events, such as the release of a compelling movie or the failure of a network link, may sometimes further increase the load on the network. As end-to-end video-over-IP streams, including high-definition channels and on-demand traffic, in conjunction with broadcast service grow as expected, CAC becomes vital to avoid oversubscription that can lead to degradation of the user experience.

**Q. How does the Cisco Integrated Video Admission Control solution for VoD and broadcast TV work?**

**A.** For more details, please see the white paper at: http://www.cisco.com/en/US/products/hw/routers/ps368/prod_white_papers_list.html.

## SERVICE ASSURANCE

**Q. What is service assurance?**

**A.** Service assurance is the detection, resiliency, and monitoring capabilities that are needed to provide service availability, service velocity, autoconfiguration of equipment, and easy end-to-end deployment through connectivity fault management and link-level protection.

**Figure 1.** Elements of Service Assurance

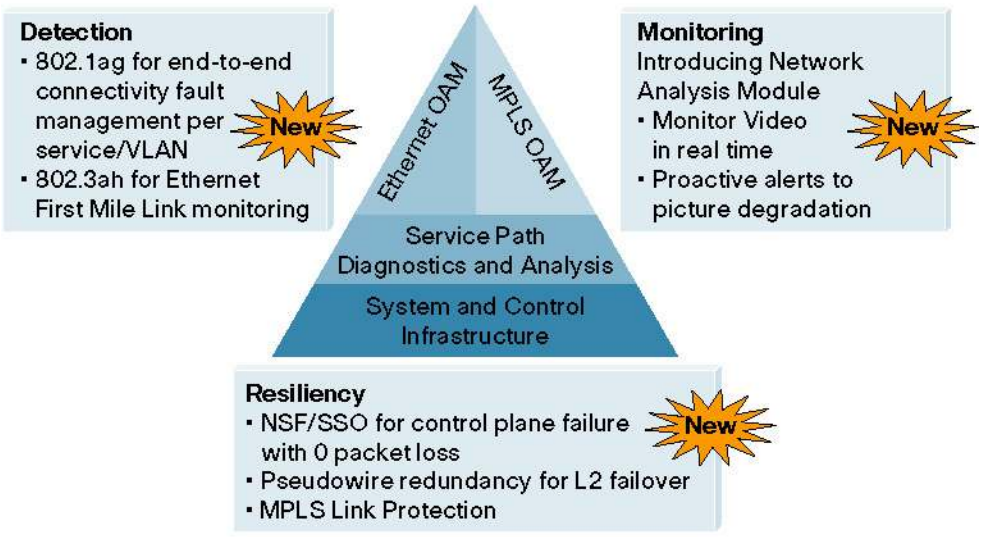


 

**Q. What does the monitoring aspect of service assurance provide?**

**A.** The monitoring aspect of service assurance provides traffic analysis integrated in the network, real-time and historical monitoring, and application response-time monitoring to determine the user experience of the network for applications such as VoD and IPTV. It therefore eases deployment, management, and support while quickly identifying the traffic with the greatest impact to performance. It also helps pinpoint latency issues to the server or to the network and helps identify problems before they impact users.

**Q. What hardware enables these monitoring capabilities?**

**A.** The Cisco Catalyst® 6500 Series/Cisco 7600 Series Network Analysis Module (NAM) brings these capabilities to the network. For more details, please visit: http://www.cisco.com/en/US/products/hw/modules/ps2706/ps5025/index.html.

## SERVICE ASSURANCE DETECTION (ETHERNET OAM)

**Q. What is the detection aspect of service assurance?**

**A.** The detection aspect refers to the Ethernet operations, administration, and maintenance (OAM), which is a general term for the management capabilities associated with Ethernet technology. In the context of Ethernet networks, it refers to the tools and utilities available to install, monitor, and troubleshoot the network.

**Q. What does the resiliency aspect of service assurance provide?**

**A.** The resiliency aspect of service assurance provides a highly available infrastructure layer with an associated set of control and signaling protocols to provide end-to-end services. It encompasses technologies and mechanisms such as MPLS Nonstop Forwarding/Stateful Switchover (NSF/SSO) for Label Distribution Protocol (LDP), Layer 3 VPN, Traffic Engineering, Fast Reroute (FRR), and Resource Reservation Protocol (RSVP), as well as mechanisms such as graceful restart capabilities for RSVP, LDP, Open Shortest Path First (OSPF), and Intermediate System-to-Intermediate System (IS-IS) Protocol.

**Q. Why is Ethernet OAM important to service providers?**

**A.** Service providers can use Ethernet OAM to extend their operational reach beyond the central office, with tools to more efficiently manage network operations. The monitoring aspect of Ethernet OAM increases visibility into the network, helping ensure everything is working correctly and knowing when it is not. Ethernet OAM raises warnings and alarms whenever a failure or degradation is detected. It also provides diagnostic and troubleshooting utilities to fix problems when they occur.

Although the use of traditional Ethernet technologies can enable carriers to save up to 85 percent in capital expenditures, because of its enterprise roots, Ethernet lacks any management capabilities inherent in the link layer. If it cannot be managed in a large carrier access network, it will not be widely deployed, despite the capital cost advantages it offers.

**Q. What is the impact of not having Ethernet OAM available?**

**A.** Without Ethernet OAM, operators lack visibility from inside the network operations center (NOC). Their only option is to send technicians into the field. Onsite visits significantly increase the operating expenses of an access network. Not only are they expensive, but they also can require a great deal of time to diagnose and troubleshoot.

**Q. What is the scope of Ethernet OAM?**

**A.** Carrier networks are typically hierarchical networks that consist of a wide variety of technologies and can span large geographical areas. Therefore, delivery of an Ethernet service could potentially span native Ethernet, Ethernet-over-SONET, Ethernet-over-MPLS, etc. transport. Thus, it is important to view Ethernet OAM as a set of layered behavior that permits monitoring and troubleshooting at a single level.

**Q. What are the layers that Ethernet OAM is required to monitor in support of Ethernet services?**

**A.** Ethernet OAM must provide capabilities on a link level and an end-to-end service level. The Link OAM provides management and troubleshooting of a single link between two Ethernet interfaces, which may be interconnected via Ethernet or an emulated Ethernet

 

connection such as Ethernet-over-SONET or Ethernet-over-MPLS. The OAM provides management and troubleshooting in a multi-hop Ethernet network for individual customer service instances.

**Q. What are the various standard bodies focusing on Ethernet OAM?**

**A.** A number of standard bodies are engaged in Ethernet OAM efforts:

- ITU-T SG 13 and SG 15
  - Ethernet OAM functionality (Y.ethoam SG 13)
  - Requirements for OAM functions in Ethernet-based networks (Y.1730 – SG 13)
- IEEE
  - 802.3ah: Ethernet in the first mile (Physical OAM)
  - 802.1ag: Connectivity management
- Metro Ethernet Forum (MEF)
  - Tracking the standards work of the other bodies, engaged in Service OAM development
- IETF
  - Tracking the standards work of the other bodies, engaged in SNMP MIB development

**Q. What is Cisco strategy for Ethernet OAM?**

**A.** Cisco is committed to the creation, development, and deployment of standards-based Link and Service OAM capabilities for Ethernet that are transparent to the underlying transport technology. Cisco is taking an active role in the standards bodies (ITU, IEEE, IETF, and MEF) responsible for defining these standards and has developed, or is in the process of developing, a number of Ethernet OAM capabilities.

## ETHERNET OAM TECHNICAL DETAILS

**Q. What is the position of IEEE 802.3ah Ethernet OAM in the OSI Reference Model?**

**A.** IEEE 802.3ah Ethernet OAM is an optional sublayer implemented in the data-link layer, between the LLC and MAC sublayers (Figure 2).

**Figure 2.** Ethernet OAM in the OSI Model

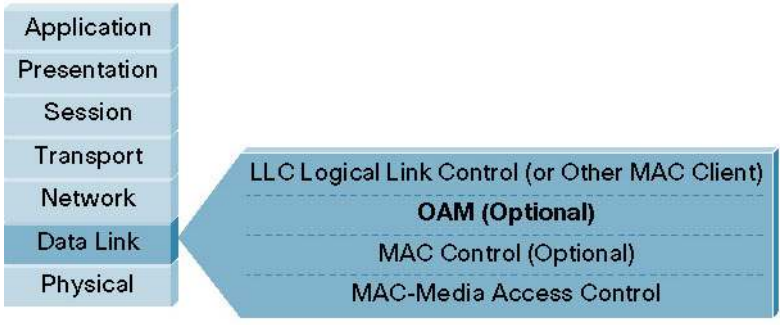


Ethernet OAM can be implemented on any full-duplex point-to-point or emulated point-to-point Ethernet link. It can be deployed for part of a system (only on particular interfaces) and does not need to be implemented systemwide.

 

**Q. What are the IEEE 802.3ah Ethernet OAM principles of operation?**

**A.** Ethernet OAM employs the following principles and concepts:

- The OAM sublayer presents a standard IEEE 802.3 MAC service interface to the superior sublayer. Superior sublayers include MAC client and link aggregation.
- The OAM sublayer employs a standard IEEE 802.3 MAC service interface to the subordinate sublayer. Subordinate sublayers include MAC and MAC control.
- Frames from superior sublayers are multiplexed within the OAM sublayer with OAM Protocol Data Units (OAMPDUs).
- The OAM sublayer parses received frames and passes OAMPDUs to the OAM client. In general, non-OAMPDUs are passed to the superior sublayer. When in OAM remote loopback mode, non-OAMPDUs are looped back to the subordinate sublayer. When the peer OAM entity is in OAM remote loopback mode, non-OAMPDUs are discarded by the OAM sublayer so that higher-layer functions (such as bridging) do not process the looped-back frames.
- Knowledge of the underlying physical-layer device is not required by the OAM sublayer.
- OAMPDUs traverse a single link and are passed between OAM client entities or OAM sublayer entities. OAMPDUs are not forwarded by OAM clients.
- OAM is extensible through the use of an organization-specific OAMPDU, organization-specific Information Type Length Value (TLV), and organization-specific Event TLV. These can be used for functions outside the scope of this standard.

**Q. Does 802.1ag Connectivity Fault Management (CFM) require all the devices in an EVC to support the protocol?**

**A.** Yes. CFM uses special, reserved MAC addresses that will not be recognized by Ethernet interfaces that do not support CFM.

**Q. What are the OAM client and sublayer?**

**A.** The OAM client is responsible for establishing and managing Ethernet OAM on a link. It also enables and configures the OAM sublayer. During the OAM Discovery phase, the OAM client monitors received OAMPDUs from the remote peer and enables OAM functionality on the link based on local and remote state as well as configuration settings. The OAM sublayer presents two standard IEEE 802.3 MAC service interfaces: one facing northbound toward the superior sublayers that include MAC client and link aggregation, and the other interface facing southbound toward the subordinate sublayer (MAC control). Furthermore, the OAM sublayer provides a dedicated interface for the OAM client to pass OAM control information and OAMPDUs to and from that client. The OAM sublayer has three constituent blocks: control block, multiplexer, and packet parser.

**Q. What is an OAMPDU?**

**A.** Ethernet OAM messages are referred to as OAM Protocol Data Units, or OAMPDUs. These are standard length, untagged, Ethernet frames within the normal frame-length bounds of 64 to 1518 bytes. The maximum OAMPDU frame size exchanged between two peers is negotiated during the discovery phase.

**Q. What is the relationship between OAM layers?**

**A.** Figure 3 shows the relation between an Ethernet virtual connection (an Ether trail) and the underlying transport virtual connection. Each of these connecting points can be considered as an Ethernet bridge providing switching between ingress and egress traffic at the Ethernet layer. Between two Ethernet connecting points, there exists a transport trail. The transport trail itself traverses through one or more transport connecting points. Transport connecting points provide switching between ingress and egress traffic at the transport layer. Basically, a connecting point at a given network layer provides switching functionality for that layer.

 

**Figure 3.** Relationship Between OAM Layers

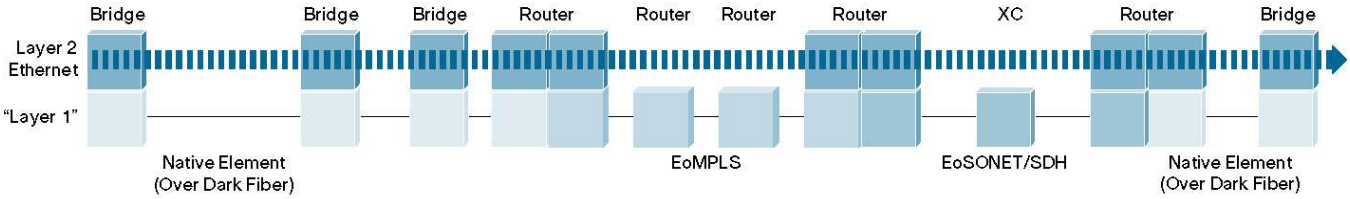


Some examples of transport trails with respect to the Ethernet client layer are SONET circuits, MPLS label-switched paths (LSPs) or pseudowires, IP tunnels (Layer 2 Tunneling Protocol Version 3 [L2TPv3]), and ATM and Frame Relay virtual circuits. A transport trail serves to connect two connecting points at the layer above it (for example, connecting two Ethernet bridges). In turn, an Ethernet trail serves to connect two connecting points at the layer above it (for example, connecting two routers). The layer that provides the service to its upper layer is known as the service layer and the upper layer that uses the service is known as the client layer.

It is important to differentiate between different network client and service layers. Each layer needs to have its own OAM mechanism. Furthermore, each layer's fault-management mechanism should operate independently from the layer below and above it. It is important to adhere to this independent layering protocol in design and development of fault-management mechanisms for Ethernet services in order to ensure ease of operation over different transport layers.

## TRANSPORT-LAYER OAM

**Q. What is Transport-Layer OAM and why is it important to service providers?**

**A.** Transport-Layer OAM provides mechanisms that are useful for monitoring link operation, such as remote fault indication and remote loopback control. In general, Transport-Layer OAM provides network operators the ability to monitor the health of the network and quickly determine the location of failing links or fault conditions. These OAM mechanisms encompass the following areas:

- Provider Edge OAM, as defined by IEEE 802.3ah
- Underlying layers of OAM, such as MPLS OAM, SONET OAM, or ATM OAM, used on various emulated Ethernet links

Given the breadth of this subject, this document will only focus on the Provider Edge OAM with regard to Transport-Layer OAM.

**Q. What is IEEE 802.3ah OAM?**

**A.** IEEE 802.3ah OAM covers the OAM frames used across a physical IEEE 802.3 medium between a provider and a customer, between two provider ports, or even potentially between two customer ports. The major objectives of 802.3ah are as follows:

- Remote failure indication
  - A mechanism is provided to indicate to a peer that the receive path of the local DTE is nonoperational. (This requires physical-layer devices that support unidirectional operation.)
  - Subscriber access physical-layer devices supporting unidirectional operation in the direction from Optical Line Termination (OLT) to Optical Network Unit (ONU) allow OAM remote failure indication from OLT during fault conditions.
  - Physical-layer devices other than those listed above do not support unidirectional operation allowing OAM remote failure indication during fault conditions.
- Remote loopback – A mechanism is provided to support a data-link-layer frame-level loopback mode.
- Link monitoring
  - A mechanism is provided to support event notification that permits the inclusion of diagnostic information.
  - A mechanism is provided to support polling of the diagnostic data.

 

- Miscellaneous
  - Implementation and activation of OAM are optional.
  - A mechanism is provided that performs OAM capability discovery.
  - An extension mechanism is provided and made available for higher-layer management applications.

## PER-SERVICE OAM

**Q. What is Service OAM and why is it important to service providers?**

**A.** Service OAM focuses on the "end-to-end" connectivity of an Ethernet virtual circuit, or Ethernet Virtual Connection (EVC). Because an Ethernet virtual circuit can be virtual, the capabilities to debug Ethernet wires such as those provided by the Transport-Layer OAM mechanisms do not necessarily extend end-to-end. IEEE 802.1ag Connectivity Fault Management uses standard Ethernet frames, which are distinguished from ordinary data frames only by destination MAC address or EtherType, which are seen and either relayed or terminated by provider bridges.

**Q. What is IEEE 802.1ag Connectivity Fault Management (CFM)?**

**A.** IEEE 802.1ag CFM provides "service" management. In other words, it allows service providers to manage each customer service instance individually. A customer service instance, or Ethernet Virtual Connection (EVC), is the service that is sold to a customer and is designated by the Service-VLAN tag. Hence, CFM operates on a per-Service-VLAN (or per-EVC) basis. It enables the service provider to know if an EVC has failed, and if so, provides the tools to enable rapid isolation of the failure.

**Q. What are 802.1ag Ethernet OAM domains?**

**A.** Domains are defined in terms of "maintenance points," which are "MACs" to IEEE 802, and "interfaces" or "ports" to others.

- A maintenance point at the edge of a domain is called a "maintenance endpoint." System administrators use maintenance endpoints to initiate and monitor CFM activity and report the results.
- A maintenance point inside a domain, and visible to a maintenance endpoint, is called a "maintenance intermediate point." Maintenance intermediate points passively receive and respond to CFM packets initiated by maintenance endpoints.

**Q. What would a domain look like?**

**A.** Figure 4 shows an example of an Ethernet OAM domain.

**Figure 4.** Ethernet OAM Domain

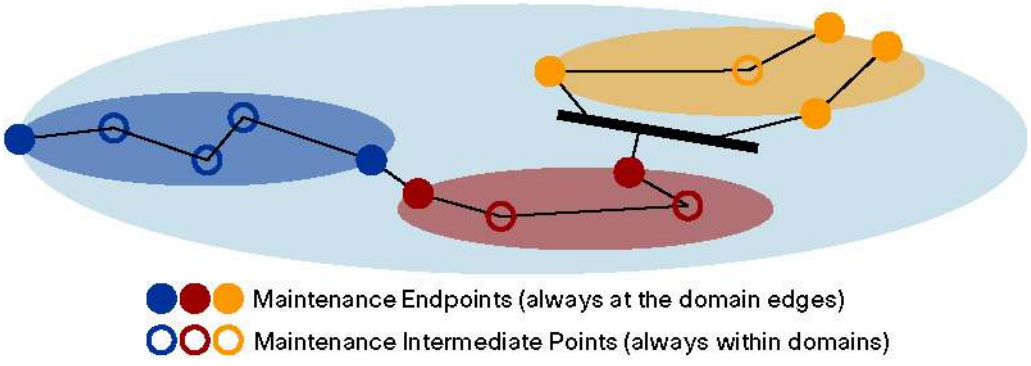


 

**Q. What are the categories of messages in 802.1ag Ethernet OAM CFM?**

**A.** Connectivity check messages, traceroute messages, loopback messages, and Alarm Indication Signal (AIS) messages.

- *Connectivity check messages* – These are "heartbeat" messages exchanged periodically between maintenance endpoints. They allow maintenance endpoints to discover other maintenance endpoints within a domain, and allow maintenance intermediate points to discover maintenance endpoints.
- *Traceroute messages* – These are transmitted by a maintenance endpoint on the request of the administrator to track the path (hop-by-hop) to a destination maintenance endpoint. They allow the transmitting node to discover vital connectivity data about the path. They are similar in concept to UDP Traceroute.
- *Loopback messages* – These are transmitted by a maintenance endpoint on the request of the administrator to verify connectivity to a particular maintenance point. Loopback indicates whether the destination is reachable or not – it does not allow hop-by-hop discovery of the path. It is similar in concept to ICMP Echo (Ping).
- *AIS messages* – These messages are generated by a maintenance endpoint or maintenance intermediate point that discovers a connectivity fault and notifies other devices of this fault.

## SERVICE-LAYER OAM

**Q. What is Service-Layer OAM and why is it important service providers?**

**A.** Service-Layer OAM provides mechanisms to support the performance monitoring of a service and also enables the customer-edge device to request and receive status from the Ethernet network so that it can configure itself to access Metro Ethernet services. The latter is an important aspect in terms of ensuring the end-to-end service-level agreement (SLA).

**Q. Are there standards for performing Service-Layer OAM?**

**A.** There is emerging work in progress within the Metro Ethernet Forum (MEF) and in the ITU SG 13 in this area. Ongoing collaboration is based on the following concepts:

- Per-service OAM mechanisms must include per-class-of-service (CoS) OAM wherein multiple classes of service can be supported within an Ethernet service, specifically for the purposes of performance management.
- The Ethernet Services OAM must support the measurement of per-service frame loss between service-flow endpoints on any two network elements supporting the same Ethernet service inside an OAM domain.
- The Ethernet Services OAM must support the measurement of two-way delay between service-flow endpoints on any network elements supporting the same Ethernet service inside an OAM domain.
- The Ethernet Services OAM should support the measurement of one-way delay between service-flow endpoints on any network elements supporting the same Ethernet service inside an OAM domain.
- The Ethernet Services OAM must support the measurement of frame delay variation between service-flow endpoints on any network elements supporting the same Ethernet service inside an OAM domain.
- The Ethernet Services OAM must support the measurement of frame delay for flooded frames between service-flow endpoints of any two network elements supporting the same Ethernet service inside an OAM domain.
- The Ethernet Services OAM must be independent of underlying transport technologies and specific transport OAM capabilities.
- The Ethernet Services OAM frames must be forwarded along on the same path as the data frames are forwarded.
- The Ethernet Service OAM should be backward compatible such that network elements that do not support OAM frames are able to forward the OAM frames in a similar fashion as the regular service/data frames.

MEF is defining the Ethernet Local Management Interface (E-LMI) to enable autoconfiguration of the CE device.

 

**Q. What are the details of the E-LMI?**

**A.** The E-LMI protocol has local significance on the UNI between and the CE device. Figure 5 shows the scope of the E-LMI protocol. (MEN – Metro Ethernet Network)

**Figure 5.** Ethernet Local Management Interface (E-LMI) Protocol

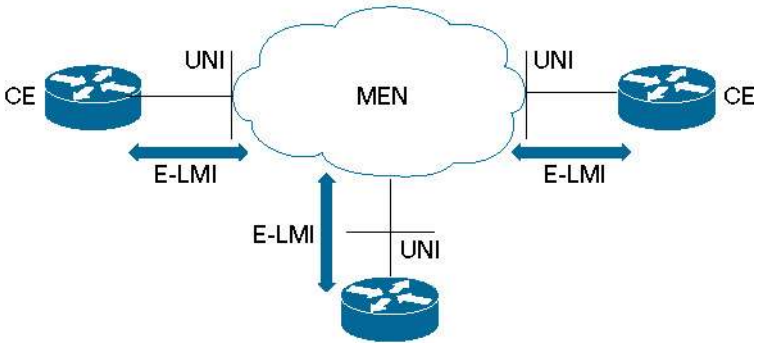


E-LMI defines the protocol and procedures that convey the information that allows autoconfiguration of the CE device. The E-LMI protocol also provides the means for notification of the status of an Ethernet Virtual Connection (EVC). In particular, the E-LMI protocol includes the following procedures:

- Notification to the CE device of the addition or deletion of an EVC
- Notification to the CE device of the availability (active) or unavailability (inactive) state of a configured EVC
- Link-integrity verification
- Communication of UNI and EVC attributes to the CE device

The E-LMI messages are transferred across the UNI using Ethernet frames.

## ETHERNET OAM-RELATED ISSUES

**Q. What are the interworking requirements for Ethernet OAM?**

**A.** Ethernet services will often be offered by service providers in combination with ATM, Frame Relay, and SONET access for large enterprise customers. This allows the enterprise customer to slowly adopt Ethernet access rather than doing so in a one-time migration. Service providers need to be able to test the end-to-end connectivity and, therefore, need interworking between the OAM messages of the different protocols. The most important internetworking function is for AIS, so that all segments of a customer service connection are notified of the alarm.

**Q. Why is Ethernet OAM important?**

**A.** Service providers have highlighted in many surveys and reports that one of the primary barriers to deployment for Metro Ethernet is the lack of management tools. Service providers must be able to quickly and easily provision and manage new services. They are looking for easy configuration, status of EVC and mapping, performance statistics, latency, jitter, packet delivery rate, and end-to-end error checking. Cisco is working with the industry to provide the most comprehensive solution for its customers.

**Q. What impact does Ethernet OAM have on the service provider's OSS?**

**A.** Ethernet OAM helps service providers to easily troubleshoot and provision services for their customers and deliver higher SLAs. Adding IEEE 802.3ah implementation in a service provider's OSS can achieve the following benefits:

- For troubleshooting, IEEE 802.3ah OAM enables service providers to monitor and troubleshoot a single Ethernet link. It is particularly valuable in the first-mile connection to the customer demarcation, where most link issues typically occur.

 

- For provisioning, 802.3ah enables the service provider to monitor a link for critical events, and then, if necessary, put the remote device into “loopback” mode to do testing on the link. It also discovers unidirectional links, which occur when only one direction of a transmission fails. Current management protocols for Ethernet do not provide the physical, link-level management enabled by 802.3ah.

## WIRELESS SERVICES MODULE (WiSM) Q & A

**Q. What is the Cisco 7600 Wireless Services Module (WiSM)?**

**A.** The Cisco 7600 Series WiSM provides unparalleled security, mobility, redundancy, and ease of use for mission-critical wireless LANs and mesh wireless LANs. It smoothly integrates into existing wireless networks based on Cisco 7600 Series Routers and Cisco Catalyst 6500 Series Switches. The module scales to deliver secure, metro wireless access to main, branch, and remote campuses. It gives operators the control they need to scale and manage their wireless networks as easily as they scale and manage their traditional wired networks. With this module, customers can realize significant total cost of ownership (TCO) benefits by reducing support costs as well as planned and unplanned network downtime. The module supports clustering capabilities of up to 3600 lightweight access points per roaming domain. It scales to 300 lightweight access points per module with support for more than 10,000 wireless client devices. For further information on this module, refer to: http://www.cisco.com/en/US/prod/collateral/modules/ps2706/ps6526/prod_qas0900aecd8036434e.shtml.

 


**Corporate Headquarters**
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

**European Headquarters**
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

**Americas Headquarters**
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

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


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

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe




Printed in USA C67-350533-00 05/06