

Ethernet OAM Technology White Paper

Issue 2.0
Date 2012-10-30

Copyright © Huawei Technologies Co., Ltd. 2012. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: <http://enterprise.huawei.com>

Contents

| | |
|---|----------|
| 1 Ethernet OAM Overview | 1 |
| 2 IEEE 802.3ah (EFM) | 1 |
| 2.1 IEEE 802.3ah Principles | 1 |
| 2.1.1 IEEE 802.3ah OAM Sublayer..... | 1 |
| 2.1.2 IEEE 802.3ah Frame Structure..... | 2 |
| 2.1.3 Operation Mode | 4 |
| 2.2 OAM Capability Discovery | 5 |
| 2.3 OAM Link Monitoring..... | 6 |
| 2.4 Fault Notification | 7 |
| 2.5 OAM Remote Loopback | 8 |
| 2.6 Remote Variable Retrieval..... | 9 |
| 2.7 Organization Specific OAMPDU..... | 10 |
| 3 IEEE 802.1ag (CFM) | 1 |
| 3.1 Multi-Domain Network Model | 1 |
| 3.2 IEEE 802.1ag Protocol Stack | 2 |
| 3.3 MD and MA | 2 |
| 3.4 MD Level | 4 |
| 3.5 MEP and MIP..... | 5 |
| 3.6 Five Functions..... | 6 |
| 3.6.1 Ethernet Fault Detection (CCM)..... | 6 |
| 3.6.2 Fault Confirmation Through Loopback | 6 |
| 3.6.3 Fault Location and Isolation Through LTM/LTR Messages | 7 |
| 3.6.4 Remote Fault Notification Through RDI | 7 |
| 3.6.5 Unexpected MEP Condition..... | 7 |
| 3.7 Format of IEEE 802.1ag Frames | 8 |
| 3.7.1 Format of Common CFM Frame Header | 8 |
| 3.7.2 Format of CCMs | 9 |
| 3.8 Differences between IEEE 802.3ah and IEEE 802.1ag..... | 10 |
| 4 ITU-T Y.1731 | 1 |
| 4.1 Basic Concepts | 1 |
| 4.1.1 ME | 1 |
| 4.1.2 MEP | 2 |

| | |
|--|----------|
| 4.1.3 MIP | 2 |
| 4.1.4 MEG Level..... | 3 |
| 4.2 Format of ITU-T Y.1731 Frames..... | 4 |
| 4.3 OAM Error Management Functions..... | 5 |
| 4.3.1 ETH-CC | 5 |
| 4.3.2 ETH-LB | 6 |
| 4.3.3 ETH-LT..... | 7 |
| 4.3.4 ETH-AIS | 7 |
| 4.3.5 ETH-RDI | 8 |
| 4.3.6 ETH-LCK | 9 |
| 4.3.7 ETH-Test..... | 9 |
| 4.4 OAM Performance Monitoring | 9 |
| 4.4.1 ETH-LM | 10 |
| 4.4.2 ETH-DM..... | 10 |
| 4.4.3 FDV | 11 |
| 4.4.4 Throughput Measurement | 11 |
| 4.4.5 Usability..... | 11 |
| 4.5 Comparison Between ITU-T Y.1731 and IEEE 802.1ag..... | 12 |
| 5 Fault Association of Ethernet OAM | 1 |
| 5.1 Association Between EFM OAM and an Interface | 1 |
| 5.2 Association Between EFM OAM and BFD | 1 |
| 5.3 Association Between EFM OAM and MPLS OAM | 1 |
| 5.4 Association Between Ethernet CFM and EFM OAM | 1 |
| 5.5 Association Between Ethernet CFM and an Interface..... | 2 |
| 5.6 Association Between Ethernet CFM and Ethernet CFM..... | 2 |
| 5.7 Association Between Ethernet CFM and BFD | 2 |
| 5.8 Association Between Ethernet CFM and MPLS OAM | 2 |
| 5.9 Association Between Ethernet CFM and Smart Link..... | 2 |
| 6 Ethernet OAM Application Scenarios | 1 |
| 6.1 Combination of EFM, CFM, and Smart Link | 1 |
| 6.1.1 Introduction..... | 1 |
| 6.1.2 Networking | 2 |
| 6.2 Combination of EFM/CFM, MPLS OAM, and BFD..... | 2 |
| 6.2.1 Introduction..... | 2 |
| 6.2.2 Networking | 3 |
| A Abbreviation | 4 |

1 Ethernet OAM Overview

To test continuity of the Ethernet Virtual Connection (EVC) at the Ethernet layer, effectively detect and locate internal network faults at the Ethernet layer, and measure usage and performance of Ethernet, an OAM mechanism independent of a client layer or service layer needs to be incorporated at the Ethernet layer to provide services that comply with the Service Level Agreement (SLA) signed with users. This requirement is crucial to the independent development of carrier-class Ethernet.

OAM capabilities of carrier-class Ethernet must meet the following requirements:

- Ethernet OAM capabilities must be independent of a client-layer or service-layer network.
- Fault management must be supported. A fault can be detected, analyzed, located, and reported to the network management system (NMS). Proper actions are then taken to rectify the fault.
- Automatic discovery and configuration management are supported. OAM capabilities must be easy to configure so that they can be widely used. The OAM capabilities can even be used on large-scale networks.
- Performance management must be supported. EVC validity and network performance parameters such as the packet loss ratio, delay, and jitter can be measured.
- OAM capabilities must be reliable, even in case of link deterioration. Therefore, the bit error correction and detection mechanism of OAM packets must be provided.
- Domain-based OAM can be provided for carriers, service providers, and users.

Standardization Progress and Current Situation

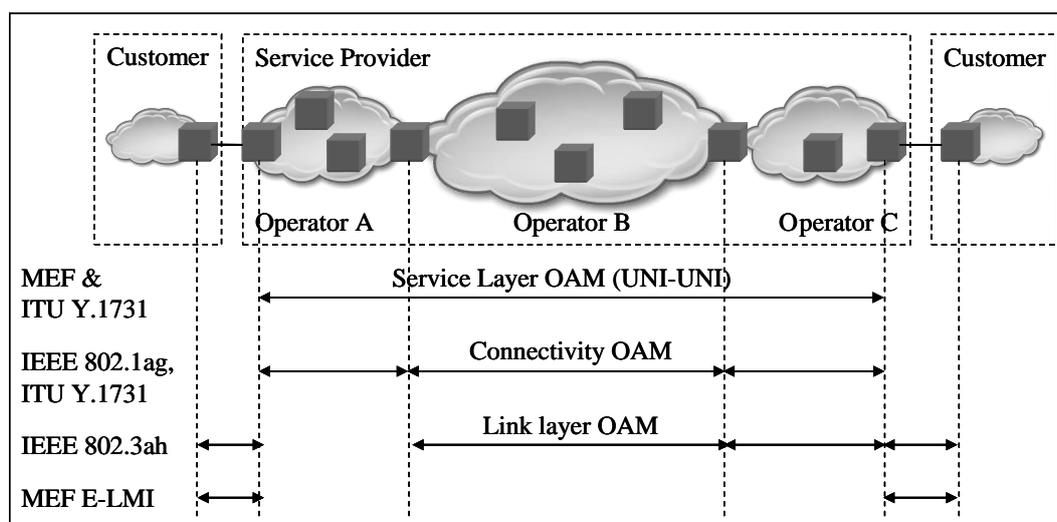
To provide services at the same level as those provided by traditional carrier-class transport networks, research communities and standardization organizations have been actively taking part in technical research and defining standards. Currently, most standards to define are related to fault management and performance management. The current progresses of these standards are as follows:

- The IEEE802 working group has constituted the following protocols:
 - IEEE 802.1ag, *Connectivity Fault Management*
 - IEEE 802.3ah, *Ethernet in the First Mile*
 - IEEE 802.1AB, *Station and Media Access Control Connectivity and Discovery*
 - IEEE 802.1ap, *Management Information Base (MIB) Definitions for VLAN Bridges*

- The IEEE 802.3ah standard has been integrated into the IEEE 802.3-2005 standard, which occurs in Chapter 57 of the IEEE 802.3-2005 standard.
- The ITU has constituted the following protocols:
 - ITU-TSG13Y.1730, *Requirements for OAM Functions in Ethernet-Based Networks*
 - ITU-TSG13Y.1731, *OAM Functions and Mechanisms for Ethernet-Based Networks*
 - ITU-T SG 15 G.8031/Y.1342, *Ethernet Protection Switching*
 - G.8032, *Ethernet Ring Protection Switching*
- The MEF has constituted the following protocols:
 - MEF7, *EMS-NMS Information Model*
 - MEF15, *Requirements for Management of Metro Ethernet Phase 1 Network Elements*
 - MEF16, *Ethernet Local Management Interface (E-LMI)*
 - MEF17, *Service OAM Requirements & Framework*
 - The *Service OAM Performance Monitoring Implementation Agreement* is being constituted.

These protocols and standards complement each other to provide end-to-end service operation management and maintenance capabilities. Figure 1-1 shows hierarchy of the OAM functions on a carrier-class Ethernet.

Figure 1-1 Hierarchy of OAM functions on a carrier-class Ethernet



Huawei switches support IEEE 802.3ah OAM and IEEE 802.1ag OAM protocols. In the two protocols, performance evaluation methods defined in the ITU-T Y.1731 protocol are referenced in calculation of performance indexes such as packet loss ratio and jitter. The major standards are described in the following chapters.

2 IEEE 802.3ah (EFM)

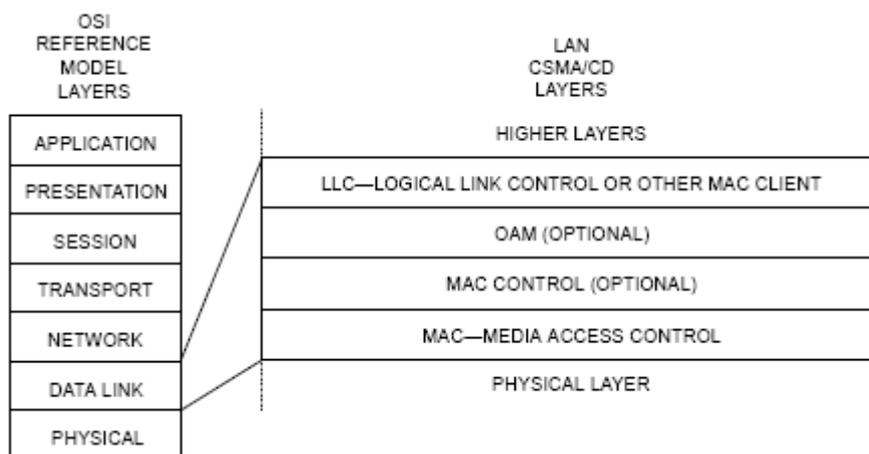
The formal IEEE 802.3ah standard, *Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks*, was released in 2004. IEEE 802.3ah provides standardized OAM suggestions and applies to direct port interconnections (in the first mile). OAM works at the link level and cannot be used for the functions that are irrelevant to single links, such as the functions of node position management, protection switching, and bandwidth reservation and allocation.

The IEEE 802.3ah defines Ethernet OAM protocol data units (OAMPDUs) to implement the OAM functions such as automatic discovery of OAM capabilities, OAM link monitoring, remote fault notification, OAM remote loopback, remote MIB obtaining, and function customization.

2.1 IEEE 802.3ah Principles

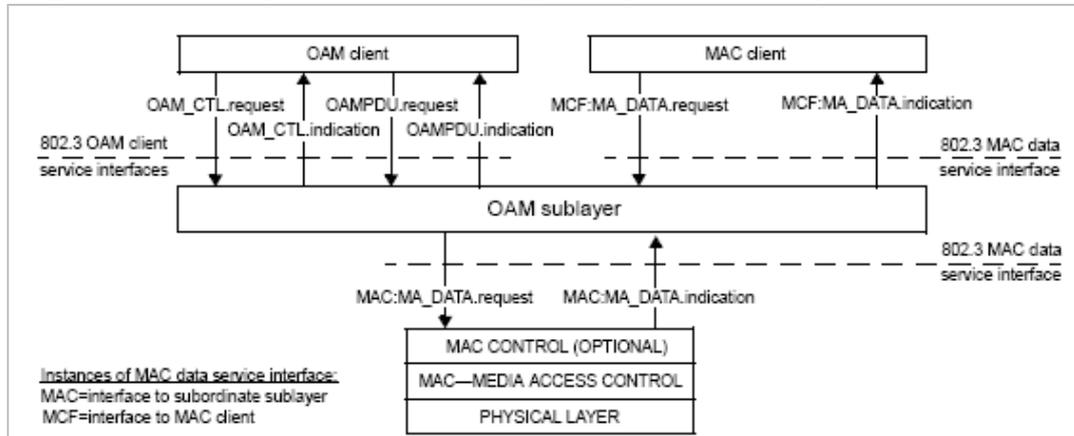
2.1.1 IEEE 802.3ah OAM Sublayer

Figure 2-1 Positions of the OAM sublayer in the OSI reference model and IEEE 802.3 CSMA/CD LAN model



The IEEE 802.3ah OAM sublayer is located at the data link layer of the OSI reference model and between the media access control (MAC) layer and the logical link control (LLC) sublayer. The IEEE 802.3ah OAM sublayer is optional. Figure 2-1 shows the positions of the OAM sublayer in the OSI reference model and IEEE 802.3 CSMA/CD LAN model.

Figure 2-2 Inter-sublayer service interfaces supported by the OAM sublayer



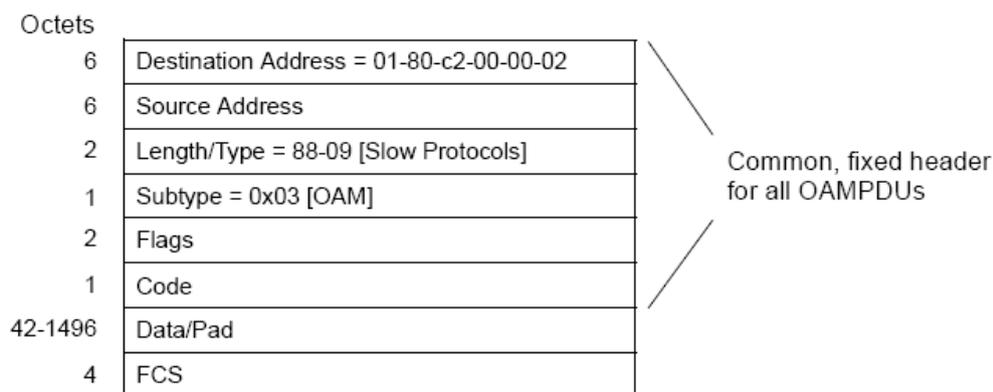
As shown in Figure 2-2, the functions of the OAM sublayer are as follows:

1. The OAM sublayer provides a standard IEEE 802.3 MAC service interface for upper-level sublayers: MAC client and link aggregation sublayers.
2. The OAM sublayer provides a standard IEEE 802.3 MAC service interface for lower-level sublayers: MAC and MAC control sublayers.
3. The OAM sublayer sorts the received OAMPDUs and transfers them to the OAM client sublayer, which is the OAM management entity. For non-OAMPDUs, the processing is as follows:
 - In normal cases, the non-OAMPDUs are transferred to the upper-level sublayers.
 - When the local devices work in remote loopback mode, the non-OAMPDUs are looped back to the lower-level sublayers.
 - When the peer OAM entity works in remote loopback mode, the non-OAMPDUs are discarded by the OAM sublayer, which prevents higher-layer functional modules, such as the bridging module, from processing loopback frames.
4. Information about the physical-layer devices is not required.
5. The OAMPDUs are transferred on a single link between two OAM clients or two OAM sublayers. The OAM client does not forward the OAMPDUs.
6. OAM can be extended through the OAMPDUs customized by organizations, such as the customized information TLVs and customized event TLVs.

The peer layers communicate with each other through the Ethernet OAMPDUs. The OAM client is the OAM management entity.

2.1.2 IEEE 802.3ah Frame Structure

The OAMPDU defined in the IEEE 802.3ah is a basic Ethernet frame, which is not a tagged frame. Figure 2-3 shows the format of an OAMPDU.

Figure 2-3 Format of an OAMPDU

Destination Address: indicates the destination address of OAM packets, which is a slow protocol multicast address. The value is 01-80-c2-00-00-02.

Source Address: indicates the MAC address of the source that sends the OAMPDUs, which is a unicast address.

Length/Type: indicates a protocol type with the length of 2 bytes. The value is 0x8809. The OAMPDU is of the slow protocol address type. The minimum frequency of sending packets of periodical fault detection is 1 second. Link Aggregation Control Protocol Data Units (LACPDUs) are of the slow protocol address type.

Subtype: indicates a subtype with the length of 1 byte. The value 3 is defined for OAM. Packets of this protocol type cannot be forwarded by bridges regardless of whether OAM capabilities are enabled or activated. The OAM packets cannot be forwarded across multiple devices.

Flags: indicates a flag that identifies an OAM event, which has the length of 2 bytes. Four types of OAM status events and three types of critical link events exist. Table 2-1 describes the three types of critical link events.

Table 2-1 Critical link events

| Critical Link Event | Description |
|---------------------|---|
| Link fault | The PHY has determined that a fault has occurred in the receive direction of the local data terminal equipment (DTE). |
| Dying gasp | An unrecoverable local failure condition has occurred. |
| Critical event | An unspecified critical event has occurred. |

Code: indicates the type of an OAMPDU, which has the length of 1 byte.

Table 2-2 OAMPDU codes

| Code | OAMPDU | Description | Source |
|-------|-----------------------|--|-------------------------|
| 00 | Information | Exchanges local and remote OAM information. | OAM client/OAM sublayer |
| 01 | Event Notification | Alerts remote DTE of link events. | OAM client |
| 02 | Variable Request | Requests one or more specific MIB variables. | OAM client |
| 03 | Variable Response | Returns one or more specific MIB variables. | OAM client |
| 04 | Loopback Control | Enables or disables OAM remote loopback. | OAM client |
| 05-FD | Reserved | Reserved | OAM client |
| FE | Organization Specific | Reserved for Organization Specific Extensions, distinguished by an organizationally unique identifier. | OAM client |
| FF | Reserved | Reserved | OAM client |

Data/Pad: indicates the Data/Pad field, in TLV format.

FCS: indicates a frame check field with the length of 4 bytes.

Note: The bidirectional transmission mode is used for OAMPDUs. Therefore, the physical layer of the Ethernet must be bidirectional.

2.1.3 Operation Mode

The OAM sublayer supports the active and passive modes. If both modes are supported, either one can be used.

The DTE that uses the active mode initialize the discovery process by sending information OAMPDUs. After the discovery process, the active DTE can transmit any OAMPDU. When the peer DTE is in passive mode, the active DTE does not respond to the OAM remote loopback commands and variable requests sent from the passive DTE.

The DTE that uses the passive mode does not initialize the discovery process. It only responds to the peer DTE during the discovery process. The passive DTE does not send variable requests or loopback control OAMPDUs.

Table 2-3 Capabilities in active and passive modes

| Capability | Active DTE | Passive DTE |
|--|------------|-------------|
| Initiates OAM Discovery process | Yes | No |
| Reacts to OAM Discovery process initiation | Yes | Yes |
| Required to send Information OAMPDUs | Yes | Yes |

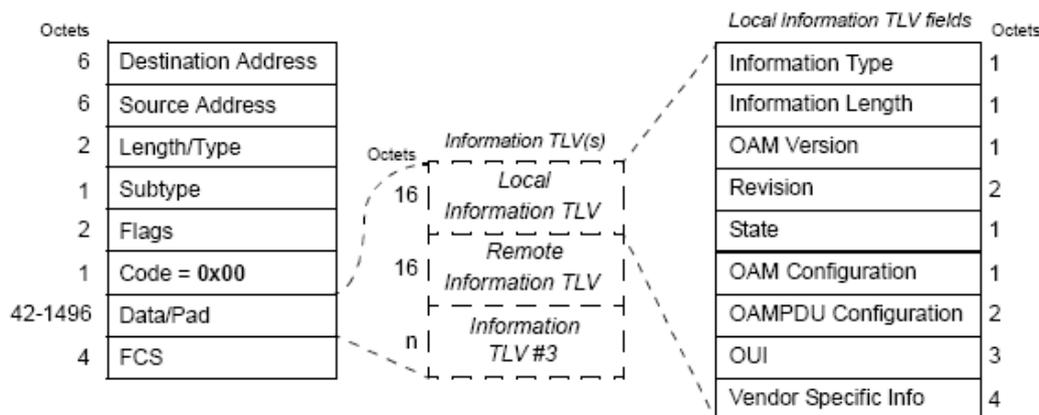
| Capability | Active DTE | Passive DTE |
|---|------------|-------------|
| Permitted to send Event Notification OAMPDUs | Yes | Yes |
| Permitted to send Variable Request OAMPDUs | Yes | No |
| Permitted to send Variable Response OAMPDUs | Yes | Yes |
| Permitted to send Loopback Control OAMPDUs | Yes | No |
| Reacts to Loopback Control OAMPDUs | Yes* | Yes |
| Permitted to send Organization Specific OAMPDUs | Yes | Yes |

* The peer DTEs must be in active mode.

2.2 OAM Capability Discovery

After Ethernet OAM capabilities are activated, a switch initiates the discovery process for the local and peer DTEs to exchange configuration and capability information. The discovery process is implemented through Information OAMPDUs. Generally, an Information OAMPDU is sent in one second. Figure 2-4 shows the format of an Information OAMPDU.

Figure 2-4 Format of an Information OAMPDU



The value 0x00 of Code indicates an Information OAMPDU used in the discovery process. The Data/Pad field uses the standard TLV structure, which includes local, remote, and customized information. Local, remote, and customized information are identified by type values. For the Data/Pad field, the type value 0x01 indicates local information, the type value 0x02 indicates remote information, and the type value 0xFE indicates customized information.

The information exchanged between the local DTE and peer DTE through the Information OAMPDUs includes the following items:

- OAM configuration information: capability information related to the configuration
- OAM mode information: active mode or passive mode

- OAMPDU information: maximum length of an OAMPDU. The minimum length is adopted before negotiation.
- Platform ID: A platform ID includes a unique Organization Unique Identifier (OUI) and the customized information of 32 bits.

Events that last more than five seconds trigger initiation of the discovery process again. These events include communications interruption of a link layer or OAMPDU packet loss.

2.3 OAM Link Monitoring

The Flags field in an OAMPDU defines three types of critical link events. For details, see 2.4 Fault Notification. The three types of critical link events cannot represent all exceptions. To understand the link status in real time and take appropriate measures for different status, the IEEE 802.3ah OAM defines another set of events that affect link operations and also defines the mechanism for notifying peer DTEs of the events. The process of the notification mechanism is as follows:

- Detecting link error events

Four types of standard fault notification events exist:

- Errored Symbol Period: indicates that the number of error signals in a unit time exceeds the defined threshold.
- Errored Frame: indicates that the number of error frames in a unit time exceeds the defined threshold.
- Errored Frame Period: indicates that the number of error frames in the specified *N* frames exceeds the defined threshold.
- Errored Frame Seconds Summary: indicates the number of seconds when error frames occur within the specified *M* seconds.

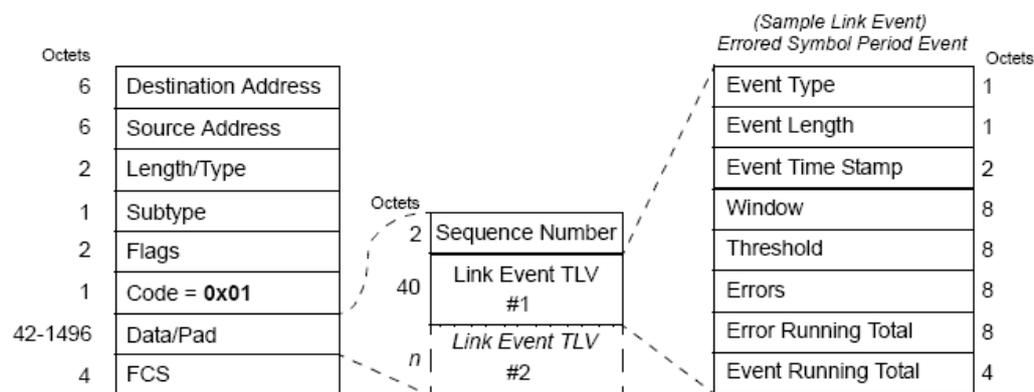
In addition, the IEEE defines customized events with the type value of 0xFE, which are optional for manufacturers.

- Sending Event Notification OAMPDU to notify the peer DTEs of detected link error events

To prevent the loss of Event Notification OAMPDU, the sender can repeatedly send Event Notification OAMPDU to ensure that the remote DTE can receive the Event Notification OAMPDU on a deteriorated link. Different events are identified by different sequence numbers.

- Querying MIB variables
- Responding to MIB variables

Figure 2-5 shows the format of an Event Notification OAMPDU.

Figure 2-5 Format of an Event Notification OAMPDU

The value 0x01 of Code indicates an Event Notification OAMPDU that is used for link monitoring. The Data/Pad field uses the standard TLV structure and four types of standard fault events exist.

The mutual notification function helps eliminate the limitation that the recipient faults can only be passively detected on traditional Ethernet. The mutual notification function enables the sender to quickly identify exceptions of the peer DTE so that faults can be fast located and rectified.

2.4 Fault Notification

The fault notification mechanism notifies the peer DTEs of severe link events on the local DTEs. When traffic is interrupted due to a fault or unavailability of a local DTE, a notification message is sent to the peer DTE through the Flags field of an OAMPDU. In IEEE 802.3ah, the following types of severe link events are defined:

Link Fault: indicates the loss of link signals on the peer DTEs.

Dying Gasp: indicates an unpredictable status such as power-down.

Critical Event: indicates an unpredictable critical event.

The switch sends a fault notification message to the remote DTE when a severe event occurs on the local DTE. The switch also logs the following fault events and reports fault events to the NMS:

- System reboot
- LPU reset
- Physical link failure
- OAMPDU timeout

The switch supports the interval for detecting received packets of 100 ms and 1s. The switch reports a timeout fault and notifies the remote DTE of the fault if no packet is received within a period five times the interval.

- Errors reported by the OAM module

When receiving a notification message, the peer logs the event contained in the message and reports it to the NMS.

 **NOTE**

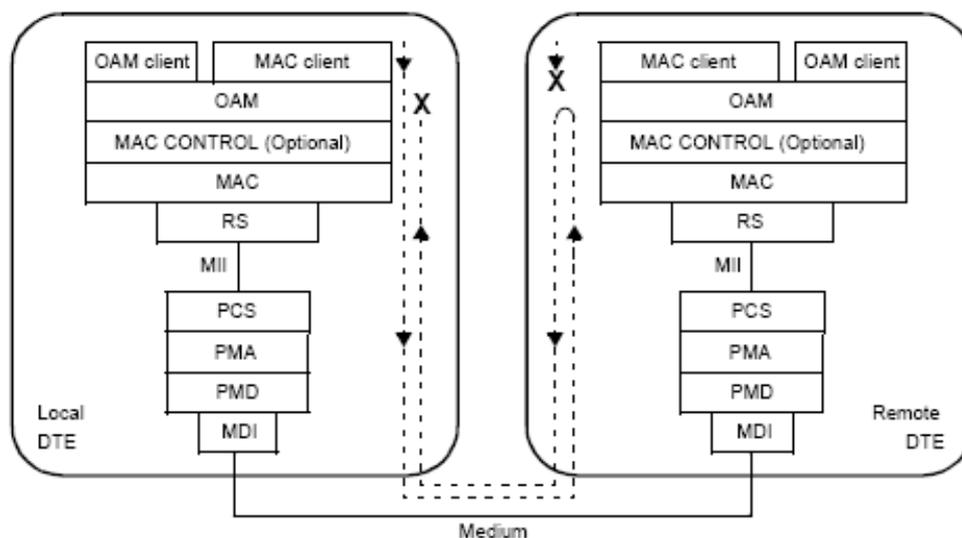
When a severe link event on the local DTE is detected, remote fault notification messages must be continuously and unidirectionally sent to the peer DTE. Messages are not sent only after the severe link event is eliminated.

2.5 OAM Remote Loopback

The OAM remote loopback function loops back frames at the data link layer, and is used for fault location and link performance test. Also, additional information about link robustness can be determined by analyzing looped-back frames. An example of this occurs when frames are lost due to a link error.

Figure 2-6 shows the path for transmitting frames in OAM remote loopback mode.

Figure 2-6 OAM remote loopback



When the local DTE sends a loopback message to the remote DTE to enable the remote DTE to enter the loopback state, OAMPDUs are received and terminated by the peer DTE. All the other non-OAMPDU frames including slow protocol frames are returned by the peer DTE without any change. If the local and peer DTEs initiate a loopback command simultaneously, the OAM compares the MAC addresses and performs the loopback operation preferentially for a lower address.

The IEEE 802.3ah OAM remote loopback can result in service interruption.

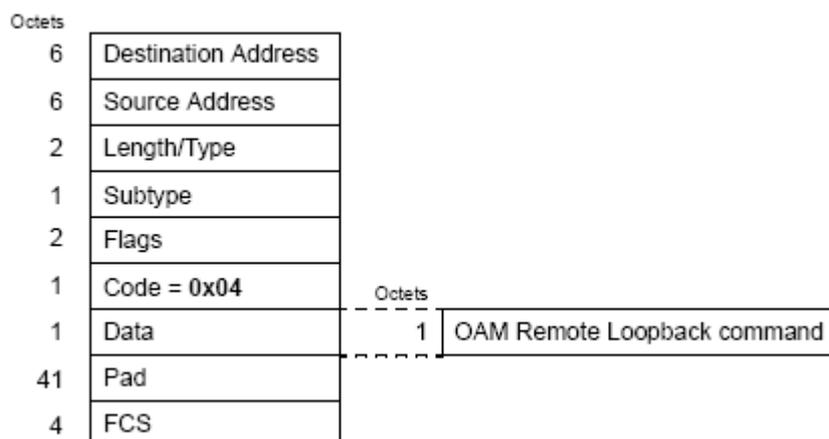
Figure 2-7 Format of a Loopback Control OAMPDU

Figure 2-7 shows the format of a remote loopback control OAMPDU. The value 0x04 of Code indicates a loopback control frame. The length of the Data field is 1 byte. Table 2-4 describes the OAM remote loopback commands.

Table 2-4 OAM remote loopback commands

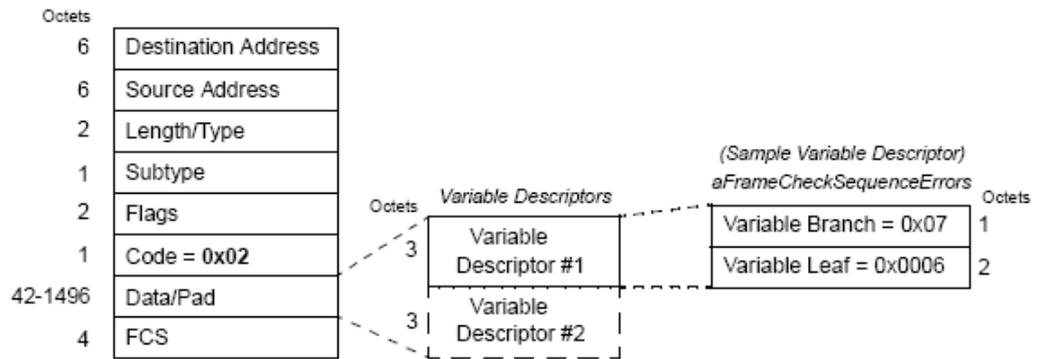
| Command | Description |
|-----------|---|
| 0x00 | Reserved - shall not be transmitted, should be ignored on reception by OAM client |
| 0x01 | Enable OAM Remote Loopback |
| 0x02 | Disable OAM Remote Loopback |
| 0x03-0xFF | Reserved - shall not be transmitted, should be ignored on reception by OAM client |

OAM remote loopback supports the test of link performance parameters such as the maximum throughput, bit error rate, delay, and jitter.

2.6 Remote Variable Retrieval

The remote variable retrieval function is used to configure remote OAM entities and obtain MIB variables to implement inband network management

Figure 2-8 Frame structure of a variable request OAMPDU

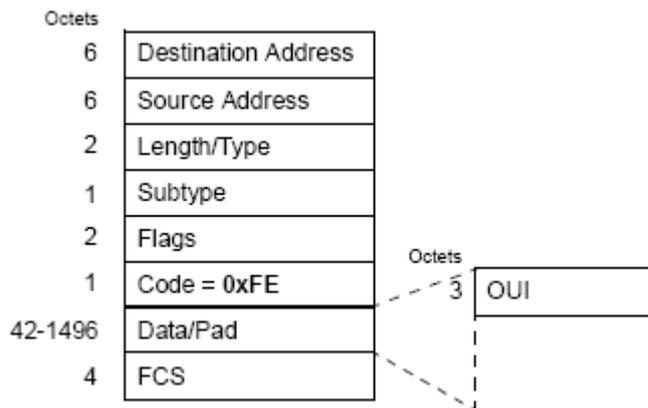


The value 0x02 of Code indicates a variable request OAMPDU. The value 0x03 of Code indicates a variable response OAMPDU.

2.7 Organization Specific OAMPDU

IEEE 802.3ah supports organization specific extensions. Device manufacturers can customize OAMPDUs, message TLVs, and event TLVs.

Figure 2-9 Format of a customized OAMPDU



As shown in Figure 2-9, the value of Code is 0xFE for a customized OAMPDU. The Data field can be defined by device manufacturers.

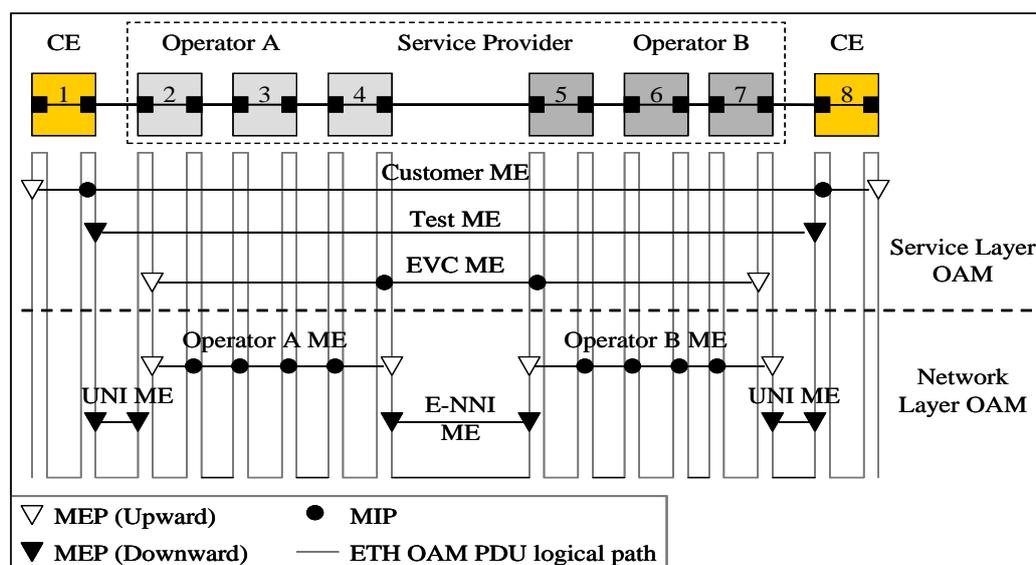
3 IEEE 802.1ag (CFM)

IEEE 802.3ah OAM applies only to point-to-point topology scenarios. For an Ethernet that needs to provide carrier-class services, the implementation of OAM between every two adjacent nodes cannot meet the requirements. IEEE 802.1ag implements an end-to-end OAM mechanism that can be applied across multiple nodes and networks.

3.1 Multi-Domain Network Model

Carrier-class Ethernet must provide different scopes and contents of management and maintenance for different entities. Generally, three roles are associated with the carrier-class Ethernet services: users such as private network users, service providers, and network carriers. A multi-domain OAM network model is uniformly used for the IEEE, ITU-T, and MEF. This multi-domain OAM network model is useful for the commercial model. Figure 3-1 shows the multi-domain OAM network model. The carrier-class Ethernet can be maintained at the user, provider, and carrier levels that correspond to different management domains. Providers are responsible for end-to-end service management, and carriers are responsible for service transmission.

Figure 3-1 Multi-domain OAM network model

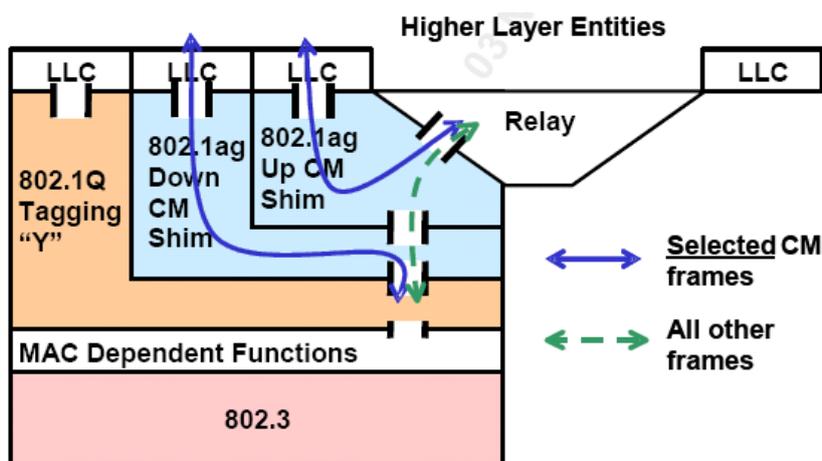


The Maintenance Association (MA) defined in the IEEE 802.1ag has the same meaning as the Maintenance Entity Group (MEG). The MA End Point corresponds to the MEP of the ITU-T, and the MA Intermediate Point corresponds to the MIP of the ITU-T.

3.2 IEEE 802.1ag Protocol Stack

IEEE 802.1ag is a link layer protocol. IEEE 802.1ag is higher than the IEEE 802.1q layer and lower than the LLC layer. Figure 3-2 shows the IEEE 802.1ag OAM protocol stack.

Figure 3-2 IEEE 802.1ag OAM protocol stack



As shown in Figure 3-2, CM refers to client management. The up shim exchanges OAM packets with the relay. The down shim (IEEE 802.1q) exchanges OAM packets with the lower layer. The LLC is used to provide an interface between the network layer and the link layer for shielding the communications modes of the lower layer, such as the Ethernet and FDDI.

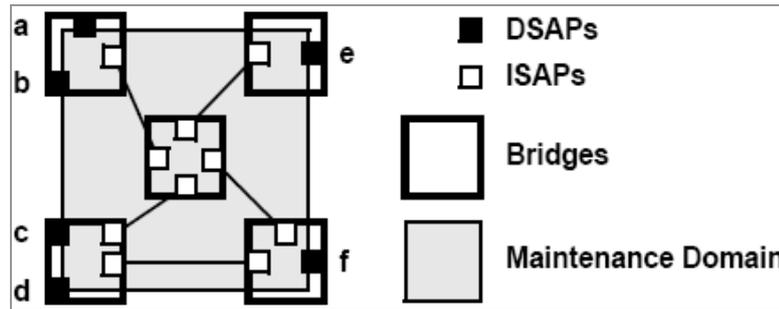
3.3 MD and MA

The maintenance domain (MD) and MA are defined in IEEE 802.1ag.

- MD

An MD can be a part or the whole of a network associated with the fault management function of IEEE 802.1ag. An MD consists of a series of domain service access points (DSAPs, which are also referred to as DoSAPs) or identifiers. A DSAP is a boundary point of the MD, such as a port of a bridge. The DSAP provides connectivity services. Intermediate service access points (ISAPs) may also exist in the MD. An ISAP is an intermediate node between two DSAPs. MDs are identified by MD names.

Figure 3-3 MD



As shown in Figure 3-3, the gray pane indicates an MD, which contains five bridges and six DSAPs. Multiple ISAPs exist in the MD. Generally, an MD is managed by an Internet service provider (ISP).

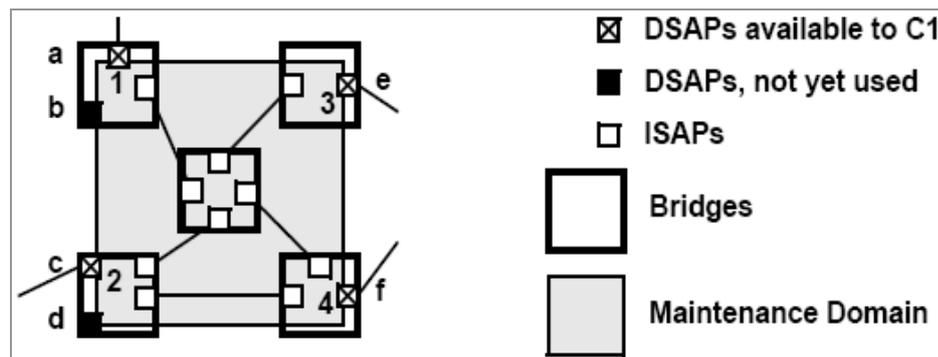
- MA

An MA is a part of an MD. Multiple service instances (SIs) can be configured in an MD. An SI can be a virtual local area network (VLAN) or a DSAP available for the services of a user. An SI consists of multiple DSAPs that are at the same MD level. An SI is identified by an ID. For details about the MD level, see 3.4 MD Level. Generally, an SI corresponds to a VLAN. An SI on a user bridge is identified by a C-VID, and an SI on a provider bridge is identified by an S-VID.

All DSAPs in an SI constitute an MA. The fault detection function is performed on all maintenance association end points (MEPs) in an MA.

A MEP is an end point of an MA. A maintenance association intermediate node is referred to as an MIP. A MEP is located at the corresponding DSAP, and an MIP is located at the corresponding ISAP. For details about the MEP and MIP, see 3.5 MEP and MIP.

Figure 3-4 MA



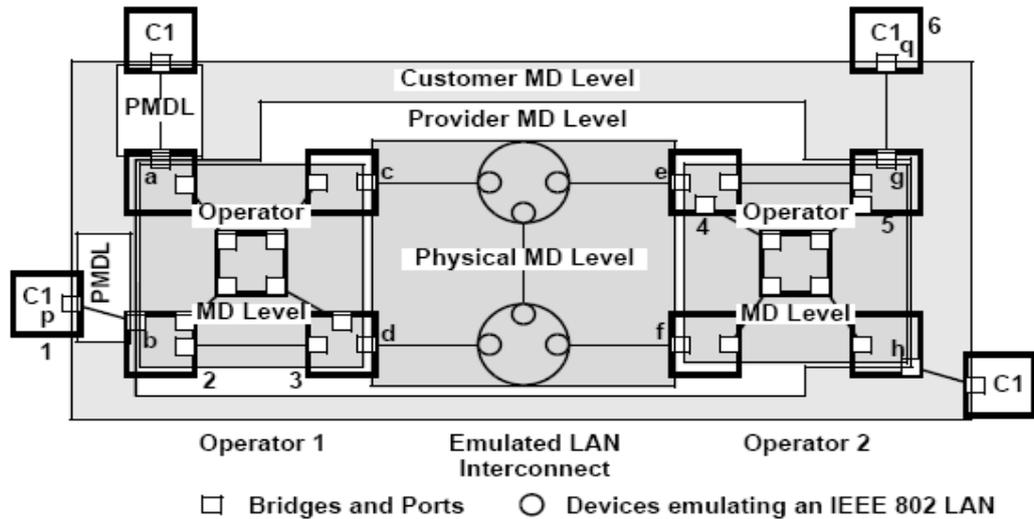
As shown in Figure 3-4, points a, c, e, and f are configured as DSAPs available for user C1. Therefore, an SI related to user C1 and the corresponding MA are created. DSAPs b and d are not used and thus do not belong to the created SI and MA.

An MA is identified by a unique MA name in the entire MD. An MA name and an MD name constitute a unique MAID. The MAID is carried in a Connectivity Fault Management (CFM) message for distinguishing an SI.

3.4 MD Level

A network can be divided into multiple MDs. These MDs can be independent of each other or overlap. The MD levels are used to indicate the nesting relationship between MDs. The value of the MD level is forwarded through a CFM PDU.

Figure 3-5 MD level



As shown in Figure 3-5, multiple overlapped MDs exist, and each MD can be divided into multiple MAs. Figure 3-5 shows seven MDs that correspond to seven SIs and seven MAs.

The provider MD level indicates SIs that are provided for user C1. Operator 1 and operator 2 create their own MAs to check integrity of the SIs provided by them. These MAs are at the operator MD level. The MD levels shown in Figure 3-5 are as follows:

Customer MD level: covered by DSAPs related to user C1. DSAPs a, b, g, and h are MIPs of the corresponding MA.

Provider MD level: formed by DSAPs a, b, g, and h. DSAPs c, d, e, and f are MIPs of the corresponding MA.

Operator MD level: Each of provider 1 and provider 2 has an MD at this level. The MD is formed by DSAPs a to h. The intermediate points between the DSAPs are MIPs of the corresponding MA.

Physical MD level: formed by DSAPs c, d, e, and f. Not all physical links require MD creation. For example, no MD is required between DSAPs g and q.

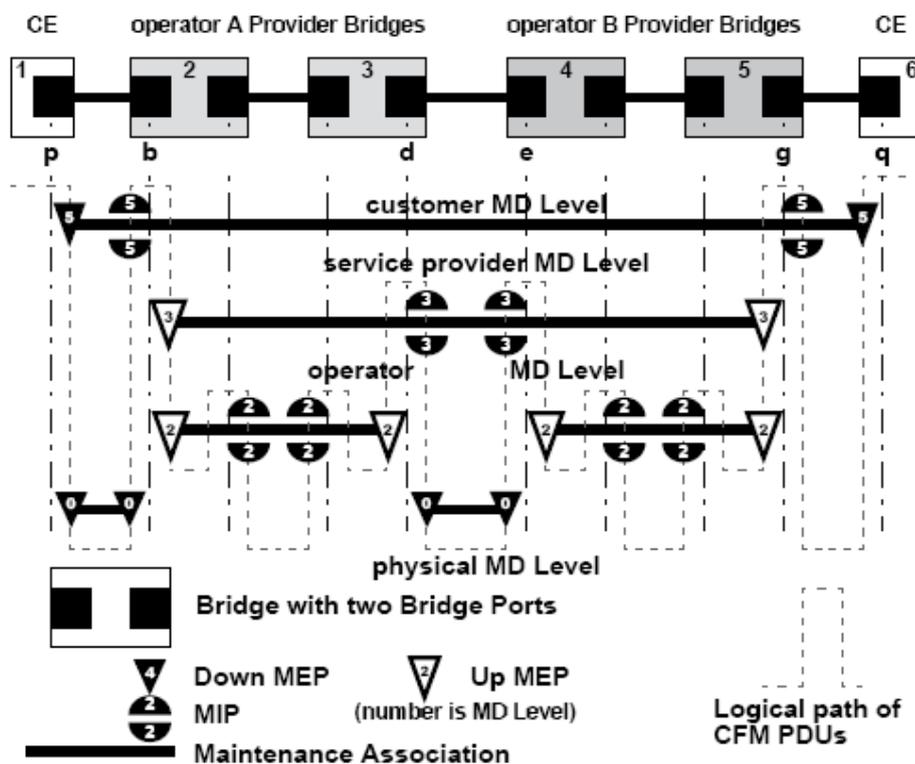
Eight MD levels from level 0 to level 7 are assigned. Level 7 indicates the highest MD level, and level 0 indicates the lowest MD level. Levels 5 to 7 are assigned for users. Levels 3 and 4 are assigned for service providers including carriers. Levels 0 to 2 are assigned for operators. Table 3-1 describes the MD levels of IEEE 802.1ag.

Table 3-1 MD levels of IEEE 802.1ag

| | | | | | | | | |
|-----------------|----------|-------------------------|---|------------------|---|----------|---|--------|
| MD Level | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
| Use | Customer | | | Service Provider | | Operator | | |
| | Highest | < - higher lower - > | | | | | | lowest |

Figure 3-6 shows hierarchy-based OAM management relationship.

Figure 3-6 Hierarchy-based OAM management relationship



An MD level indicates the corresponding service level of OAM. MDs related to an MD level can be nested but not be intersected. The principle of processing PDUs of different MD levels is transparently transmitting higher-level PDUs and blocking lower-level PDUs.

3.5 MEP and MIP

MEPs and MIPs are two types of IEEE 802.1ag OAM entities and are referred to as maintenance points (MPs).

- Maintenance Association End Point (MEP)
A MEP is formed by a set of DSAPs that belong to an SI. A MEP is an end point of an MA. For details, see 3.3 MD and MA. A MEP can generate and receive CFM PDUs and

can also respond to the received CFM PDUs. The point-to-point relationship between any two MEPs in an MA is referred to as a maintenance entity (ME).

The MEP periodically sends continuity check messages (CCMs) and discards any received CFM PDUs that have lower MD levels than the MEP. The MEP can send loopback messages (LBMs) and linktrace messages (LTMs) and can also respond to LBMs or LTMs through the loopback reply messages (LBRs) or linktrace reply messages (LTRs).

Each MEP is identified by a unique MEPID in an MA. Therefore, the combination of an MAID, an MD Level, and an MEPID can identify a MEP on a network.

- Maintenance Association Intermediate Point (MIP)

An MIP is an intermediate node between MEPs in an MA. The MIP responds to only the received CFM PDUs and does not generate any CFM PDUs. An MIP supports the following functions:

- Authenticating the received CFM PDUs
- Responding to LBMs with LBRs
- Forwarding the received LTMs and responding to the LTMs with the LTRs

3.6 Five Functions

IEEE 802.1ag implements network-level OAM. IEEE 802.1ag is used for simulating connections of bridges (VLAN-aware) on a LAN and providing functions of detecting, locating, and separating network faults. The transmission fault management of the IEEE 802.1ag supports the following functions:

3.6.1 Ethernet Fault Detection (CCM)

The fault detection function is implemented through CCMs. Each MEP periodically multicasts CCMs to all the other MEPs in an MA. The receiver can update the status of the connection with the peer MEP. If a MEP does not receive any CCMs from the peer MEP within a period 3.5 times the interval, the MEP considers that a peer fault or link fault occurs and then reports the fault to the administrators and service users. The MEP or administrator initiates fault confirmation, location, and isolation.

The interval for sending CCMs is configurable. In IEEE 802.1ag, the following intervals are defined: 3.3 ms, 10 ms, 100 ms, 1s, 10s, 1 minute, and 10 minutes.

The following faults can be detected through CCMs:

- Connectivity loss (no CCM is received within a period 3.5 times the interval)
- Configuration errors
- Hardware faults of the two ends such as a link or node fault
- Software faults of the two ends such as software errors and memory crash

CCMs and LTMs are multicasted in an MA.

3.6.2 Fault Confirmation Through Loopback

An LBM is unicast to check whether a fault occurs on the link between the local and peer DTEs. The receive node checks the LBM. If conditions are met, the receive node sends an LBR packet (a unicast MAC message) to the source node. The source node reports an alarm if

it does not receive any LBR within five seconds after sending the LBM. Intermediate nodes perform only Layer 2 forwarding.

LBMs can carry additional data that must be returned in LBRs without change. Therefore, LBMs can also be used for traffic test.

3.6.3 Fault Location and Isolation Through LTM/LTR Messages

LTMs are used for locating and isolating faults at the MAC layer on a network. An LTM contains the following fields:

DA: indicates the DA field of an Ethernet frame, which is a multicast address.

SA: indicates the SA field of an Ethernet frame, which is the address of the MIP that forwards an LTM.

Original MAC Address: indicates the MEP address that initiates an LTM.

Target MAC Address: indicates a destination MEP address.

An LTM is multicast from a MEP to another MEP. The LTM carries the MAC address of the destination node, and is sent to the protocol entities on the network hop-by-hop. After receiving the LTM, each protocol entity processes the target MAC address carried in the LTM. The protocol entity queries the forwarding table according to the target MAC address and adds the MAC address of the local DTE to the LTM. The protocol entity forwards the LTM through the corresponding egress and sends an LTR (a unicast message) to the source node. This process continues until the LTM reaches the destination that is identified by the target MAC address, the query of the forwarding table fails, or the boundary of the MA is reached. In this way, the source node can obtain the MAC addresses of all nodes on the forwarding path.

The source MEP reports an alarm if the source MEP does not receive any LTR within five seconds after sending an LTM.

3.6.4 Remote Fault Notification Through RDI

The fault notification function is performed to notify the node upstream and downstream of the fault information. Local and peer DTEs send CCMs to each other. If a DTE does not receive a CCM, the sender sends a CCM with the remote defect indication (RDI) to the DTE.

The alarm suppression function is performed to sort alarms based on alarm priorities, report only alarms of high priorities, and suppress alarms of low priorities. Through alarm suppression, network congestion and even network crash caused by a large number of notification messages on the network can be avoided. In addition, faults can be located quickly.

3.6.5 Unexpected MEP Condition

A MEP detects an unexpected MEP when it receives a CCM frame with a correct MEG Level, such as an MEG Level equal to the MEP's own MEG Level, a correct MEG ID, but an unexpected MEP ID, which includes the MEP's own MEP ID. Determining unexpected MEP IDs is possible when the MEP maintains a list of its peer MEP IDs. A list of peer MEP IDs must be configured on each MEP during provisioning. This fault is most likely caused by incorrect configuration.

Entry criteria: A MEP receives a CCM with correct MEG level, correct MEG ID, but with unexpected MEP ID.

Exit criteria: During an interval equal to 3.5 times the CCM transmission periods, the MEP does not receive CCMs with an unexpected MEP ID.

3.7 Format of IEEE 802.1ag Frames

For an Ethernet frame that carries an OAM message, the SA is the address of the MP that sends the OAM frame. The DA can indicate the address of the following types of nodes:

- All MEPs in an MA, such as CCMs
- All MPs in an MA, such as LTMs
- An MP in an MA, such as an LBM, LBR, or LTR

3.7.1 Format of Common CFM Frame Header

Figure 3-7 Format of common CFM frame header

| | Octet |
|-----------------------------|-----------------------|
| MD Level | 1 (high-order 3 bits) |
| Version | 1 (low-order 5 bits) |
| OpCode | 2 |
| Flags | 3 |
| First TLV Offset | 4 |
| Varies with value of OpCode | 5 |
| End TLV (0) | First TLV Offset + 5 |

As shown in Figure 3-7, a common CFM frame header contains the following fields:

- MD Level: maintenance level. This field has a length of 3 bits. The value ranges from 0 to 7. A larger value indicates a higher MD level.
- Version: This field has a length of 5 bits. The value is always 0.
- OpCode: message code. This field has a length of 2 bytes and indicates a unique message type. Only 1 byte is used.

Table 3-2 Value assignment of the OpCode field

| CFM PDU or Organization | OpCode Range |
|--------------------------------|--------------|
| Reserved for IEEE 802.1 | 0 |
| Continuity Check Message (CCM) | 1 |
| Loopback Reply (LBR) | 2 |
| Loopback Message (LBM) | 3 |
| Linktrace Reply (LTR) | 4 |
| Linktrace Message (LTM) | 5 |
| Reserved for IEEE 802.1 | 6 - 31 |

| CFM PDU or Organization | OpCode Range |
|-------------------------|--------------|
| Defined by ITU-T Y.1731 | 32 – 63 |
| Reserved for IEEE 802.1 | 64 – 255 |

Flag: This field has a length of 3 bytes and its value is determined by a specific OpCode.

First TLV Offset: offset address of the first TLV. This field has a length of 4 bytes.

Varies with Value of Opcode: This field has an unfixed length. This field uses the TLV format, as shown in Figure 3-8.

Figure 3-8 TLV format

| | Octet |
|---------|-------|
| Type | 1 |
| Length | 2 – 3 |
| (Value) | 4 |

IEEE 802.1ag allows any organizations to customize TLVs. The type value, however, must be 31.

End TLV: This field indicates the end of a CFM PDU. This field contains only the Type field that has a length of 1 byte and a value of 0. This field does not contain the Length and Value fields.

3.7.2 Format of CCMs

Figure 3-9 Format of CCMs

| | Octet |
|---|--|
| Common CFM Header | 1 – 4 |
| Sequence Number | 5 – 8 |
| Maintenance association End Point Identifier | 9 – 10 |
| Maintenance Association Identifier (MAID) | 11 – 58 |
| Defined by ITU-T Y.1731 | 59 – 74 |
| Reserved for definition in future versions of the protocol ^a | |
| Optional CCM TLVs | First TLV Offset + 5 ^b |
| End TLV (0) | First TLV Offset + 5, if no Optional CCM TLVs are present |

As shown in Figure 3-9, a CCM contains fields such as Sequence Number, MEPID, and MAID in addition to the Common CFM Header field. In the Flag field of a CCM header, the

highest bit 1 indicates an RDI, and the lowest bit 3 indicates a CCM interval. Other bits are reserved.

Figure 3-10 CCM interval

| Transmission Interval | max. CCM Lifetime | min. CCM Lifetime | CCM Interval field |
|-----------------------|-------------------|-------------------|--------------------|
| invalid | | | 0 |
| 3 1/3 ms (300 Hz) | 11 2/3 ms | 10 5/6 ms | 1 |
| 10 ms | 35 ms | 32.5 ms | 2 |
| 100 ms | 350 ms | 325 ms | 3 |
| 1 s | 3.5 s | 3.25 s | 4 |
| 10 s | 35 s | 32.5 s | 5 |
| 1 min | 3.5 min | 3.25 min | 6 |
| 10 min | 35 min | 32.5 min | 7 |

3.8 Differences between IEEE 802.3ah and IEEE 802.1ag

IEEE 802.1ag and IEEE 802.3ah are both Ethernet OAM standards. However, their application scenarios and implemented functions of IEEE 802.1ag and IEEE 802.3ah differ from each other. Table 3-3 describes the differences between IEEE 802.3ah and the IEEE 802.1ag.

Table 3-3 Differences between IEEE 802.3ah and the IEEE 802.1ag

| Item | IEEE 802.3ah | IEEE 802.1ag |
|--------------------|---|--|
| Physical media | Networks of the IEEE 802.3 type | All physical media networks that transmit frames in 802 format |
| Loopback | Remote loopback is implemented. Data is looped back, which affects services. | Data is not looped back. The processing is similar to that of the ping request and its response, and does not affect the normal data transmission. |
| | None. | A loopback test is performed at any rate within the rate range of a physical link. |
| Connectivity check | The connectivity check interval is 1 second. Existing connections can be determined, but unexpected connections cannot be determined. | The connectivity check interval varies. Existing and unexpected connections can be determined. |
| Monitoring range | Only a single link rather than the shared link can be monitored. | A single link or even a part or the whole of a network can be monitored. The shared link can also be monitored. |

| Item | IEEE 802.3ah | IEEE 802.1ag |
|--|--|--|
| Functions | Obtaining the statistics on a remote interface | None. |
| | Sending alarms when the number of error messages reaches the threshold | None. |
| | None. | LT function |
| Extensibility | Manufacturers are allowed to customize TLVs and use private operation codes to customize new functions. | Manufacturers are allowed to customize only the TLVs of their own and are not allowed to customize new functions. |
| Application scenarios and restrictions | IEEE 802.3ah is applicable to point-to-point OAM scenarios and is more suitable for the physical layer. IEEE 802.3ah cannot be applied across bridges. You do not need to pay attention to complex services. | IEEE 802.1ag is applicable to end-to-end OAM scenarios. Simulation of complex service scenarios is required for OAM packets. |

4 ITU-T Y.1731

ITU-T Y.1731 was defined together with IEEE 802.1ag. The initial draft of ITU-T Y.1731 was released earlier than IEEE 802.1ag. ITU-T Y.1731 and IEEE 802.1ag use the same packet format. ITU-T Y.1731 defines more comprehensive functions than IEEE 802.1ag and is a supplement to IEEE 802.1ag. Many functions defined in ITU-T Y.1731 have not been implemented by certain manufacturers. Some of the terms defined in ITU-T Y.1731 and IEEE 802.1ag are the same. Table 4-1 describes the differences between terms defined in ITU-T Y.1731 and IEEE 802.1ag.

Table 4-1 Differences between some terms defined in the ITU-T Y.1731 and IEEE 802.1ag

| IEEE Std 802.1ag-2007 | ITU-T Y.1731 (2006) | Description |
|--|------------------------|--|
| MA | MEG | |
| MAID (Domain name + short name of an MA) | MEGID | In ITU-T Y.1731, the MEGID does not consist of a domain name and the short name of an MEG, which is different from the IEEE 802.1ag. |
| MD | The MD is not defined. | |
| MD Level | MEG Level | |

4.1 Basic Concepts

4.1.1 ME

A maintenance entity (ME) is an entity that requires management and indicates the relationship between two maintenance entity group points (MEPs). For details about the MEP, see 4.1.2 MEP. Figure 4-1 shows examples of MEs on the Ethernet. MEs can be nested but cannot overlap.

Figure 4-1 ME examples

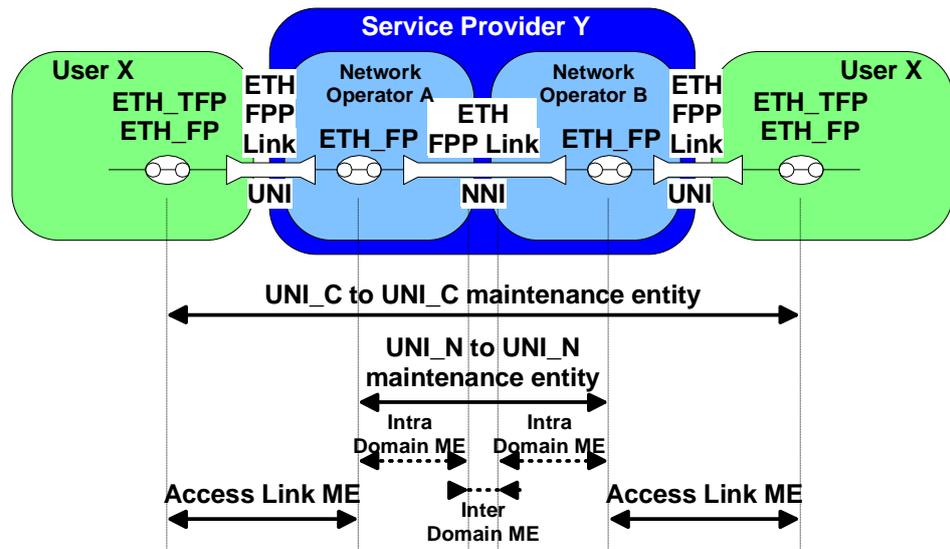


Table 4-2 Mapping between MEs

| | |
|-------------------|--|
| UNI_C to UNI-C ME | UNI-UNI (customer) |
| UNI_N to UNI_N ME | UNI-UNI (provider) |
| Intra-domain ME | Network segment inside a provider (PE-PE) |
| Inter-domain ME | Network segment between providers (PE-PE) |
| Access link ME | ETH link OAM to UNI (customer to provider) |
| Inter-domain ME | ETH link OAM to NNI (carrier to carrier) |

An ME group (MEG) includes different MEs that meet the following conditions:

- MEs in an MEG exist in the same administrative boundary.
- MEs in an MEG have the same MEG level (see 4.1.4 MEG Level).
- MEs in an MEG belong to the same point-to-point ETH connection or multipoint ETH connectivity.

4.1.2 MEP

A MEP indicates an end point of an ETH MEG. A MEP can send and terminate OAM frames for error management and performance monitoring.

A MEP can monitor the signal flow to perform functions, such as counting frames. The monitoring does not interrupt ETH service forwarding.

4.1.3 MIP

An MEG intermediate point (MIP) can respond to certain OAM frames. The MIP does not originate any OAM frames or process any forwarded ETH services.

4.1.4 MEG Level

The MEG level is similar to the MD level in IEEE 802.1ag. When MEGs are nested, the OAM flow of each MEG must be identified so that the OAM flow can be distinguished from the OAM flows of other MEGs. This procedure prevents OAM flow leakage from the corresponding OAM domains.

ITU-T Y.1731 defines eight MEG levels (levels 0 to 7) to apply to the actual network deployment scenarios. If signal flows of data channels of customers, providers, and carriers cannot be distinguished based on ETH layer encapsulation, signal flows can be assigned different MEG levels to distinguish between OAM frames. The MEG level assignment rules are as follows:

- MEG levels 5, 6, and 7 are assigned for customers.
- MEG levels 3 and 4 are assigned for providers.
- MEG levels 0, 1, and 2 can be assigned for carriers.

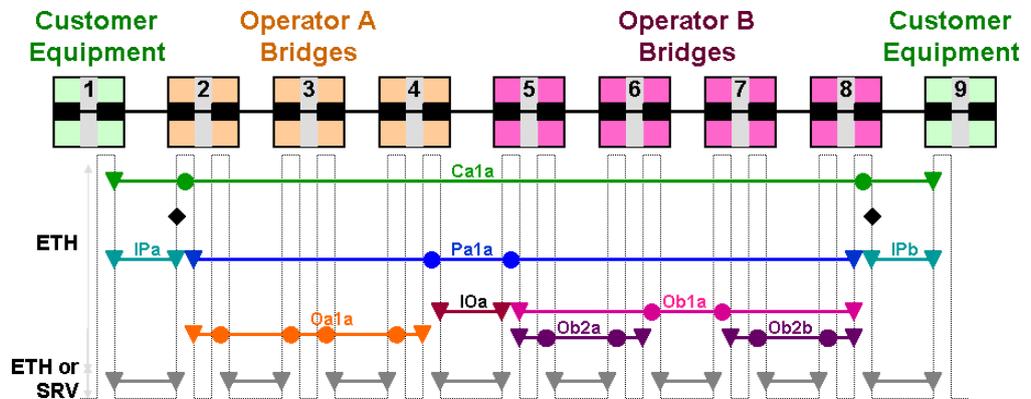
OAM frames of higher MEG levels can be transparently transmitted by MEGs of lower levels. OAM frames of lower MEG levels can be terminated by MEGs of higher levels. Table 4-3 lists the MEG levels.

Table 4-3 MEG levels

| MEG Level | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|-----------|----------|-------------------------|---|----------|---|----------|---|--------|
| Use | Customer | | | Provider | | Operator | | |
| Priority | Highest | < - Higher Lower - > | | | | | | Lowest |

Figure 4-2 shows an example of an MEG level assignment. The customer, provider, and carrier are assigned MEG levels. The default MEG level assignment is performed. As shown in Figure 4-2, triangles indicate MEPs; small cycles indicate MIPs; and diamonds indicate TrCPs.

Figure 4-2 MEG level assignment



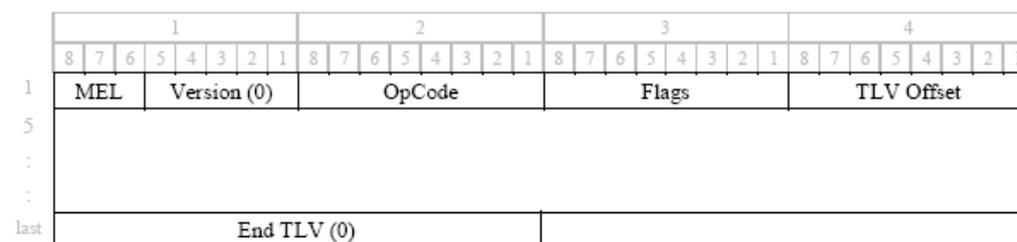
- For the UNI_C to UNI_C customer ME (Ca1a), MEG level 5 can be assigned.

- For the UNI_N to UNI_N provider ME (Pa1a), MEG level 4 can be assigned.
- For the end-to-end carrier MEs (Oa1a and Ob1a), MEG level 2 can be assigned.
- If MEs Ob2a and Ob2b are required on the network of operator B, a lower MEG level such as MEG level 1 can be assigned.
- For the UNI_C to UNI_N MEs (IPa and IPb) between the customer and the provider, MEG level 0 can be assigned.
- For the IOa between carriers, MEG level 0 can be assigned.

4.2 Format of ITU-T Y.1731 Frames

ITU-T Y.1731 OAMPDUs are carried in Ethernet frames. Some fields are shared by all OAMPDUs, as shown in Figure 4-3.

Figure 4-3 Format of shared OAM information elements



1. MEL: indicates the MEG level of an OAMPDU. This field has a length of 3 bits. The value is an integer that ranges from 0 to 7. For details about the MEG level, see 4.1.4 MEG Level.
2. Version: indicates the OAM version. This field has a length of 5 bits. The value is an integer. If the version number of an OAM frame is different from this value, the OAM is discarded.
3. OpCode: indicates the type of an OAMPDU. This field has a length of 8 bits. The **OpCode** field identifies the other fields in an OAMPDU. Table 4-4 describes the values of the **OpCode** field.

Table 4-4 Values of the OpCode field

| OpCode Value | OAMPDU Type | MEP/MIP Related to the OpCode |
|---|-------------|----------------------------------|
| OpCode values that are defined in the ITU-T Y.1731 and IEEE 802.1 | | |
| 1 | CCM | MEP |
| 3 | LBM | MEP and MIP (connectivity check) |
| 2 | LBR | MEP and MIP (connectivity check) |
| 5 | LTM | MEP and MIP |
| 4 | LTR | MEP and MIP |

| OpCode Value | OAMPDU Type | MEP/MIP Related to the OpCode |
|---|---|-------------------------------|
| 0, 6 to 31, and 64 to 255 | Reserved for use in the IEEE 802.1 | |
| OpCode values that are defined in only the ITU-T Y.1731 | | |
| 33 | AIS | MEP |
| 35 | LCK | MEP |
| 37 | TST | MEP |
| 39 | APS | MEP |
| 41 | MCC | MEP |
| 43 | LMM | MEP |
| 42 | LMR | MEP |
| 45 | IDM | MEP |
| 47 | DMM | MEP |
| 46 | DMR | MEP |
| 49 | EXM | |
| 48 | EXR | |
| 51 | VSM | |
| 50 | VSR | |
| 32, 34, 36, 38, 44, and 52 to 63 | Reserved for future ITU-T standardization | |

4. Flag: This field has a length of 8 bits. The use of each bit depends on the OAMPDU type.
5. TLV Offset: indicates the offset of the first TLV relative to the TLV Offset field in an OAMPDU. This field has a length of 8 bits. The value of this field depends on the OAMPDU type. If the value is 0, it indicates the first byte after the TLV Offset field.
6. End TLV: indicates the end of an OAMPDU. The type value is 0. The Length and Value fields are not contained.

4.3 OAM Error Management Functions

4.3.1 ETH-CC

The Ethernet connectivity check (ETH-CC) function detects the loss of continuity (LoC) between two MEPs in an MEG. ETH-CC can detect the unnecessary connection between two MEGs, improper MEP connection in an MEG, and other faults such as an incorrect MEG level and an incorrect error period. ETH-CC can be used for error detection, performance monitoring, and protection switching.

Generally, the ETH-CC function is performed between two MEPs. If a MEP does not receive any CCM from a peer MEP in the peer MEP list within a period 3.5 times the ETH-CC transmission interval, the MEP considers that the connection between the MEP and the peer MEP is torn down. The MIP transparently transmits the ETH-CC message.

The OAMPDU used for carrying an ETH-CC message is a CCM. The frame that carries a CCM PDU is a CCM frame. A CCM PDU carries the following important fields:

- MEP ID: indicates a MEP in an MEG. This field has a length of 16 bits. The highest bit 3 of the first byte is set to 0.
- MEG ID: indicates the MEG that a MEP belongs to. This field has a length of 48 bytes.
- Flag: This field contains the RDI and Period fields.

Figure 4-4 Format of the Flag field in a CCM PDU

| MSB | | | | LSB | | | |
|-----|--------------|---|---|-----|--------|---|---|
| 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 |
| RDI | Reserved (0) | | | | Period | | |

- RDI: indicates an RDI when the value of bit 8 is set to 1; otherwise, the value of this field is 0.
- CCM frame transmission interval: This field is configurable. The following values can be set:
 - 3.33 ms (default transmission interval for protection switching)
 - 10 ms
 - 100 ms (default transmission interval for performance monitoring)
 - 1s (default transmission interval for error management)
 - 10s
 - 1 minute
 - 10 minutes

4.3.2 ETH-LB

The Ethernet loopback (ETH-LB) function is performed to check connectivity between a MEP and an MIP or a peer MEP. The unicast loopback and multicast loopback are supported.

Unicast Loopback

The Ethernet unicast loopback function supports the following applications:

- This function checks bidirectional connectivity between two peer MEPs or between a MEP and an MIP.
- Bidirectional online and offline diagnosis and test can be performed between two MEPs. For example, a bandwidth traffic test or bit error rate test can be performed.
- This function uses the LBM and LBR. The RMIP or MEP replies with a unicast LBR after receiving a unicast LBM destined for the remote MIP or MEP itself.

An MIP discards a received LBR frame destined for the MIP itself because the LBR frame is invalid.

Multicast Loopback

The Ethernet multicast loopback function checks connectivity between a MEP and multiple peer MEPs. This function uses the LBM and LBR. The multicast frame that carries an LBM PDU is referred to as a multicast LBM frame.

When a MEP initiates a multicast loopback request, the multicast LBM frames that carry the specified IDs are sent to the peer MEPs in the same MEG. By default, the MEP receives unicast LBR frames from the peer MEPs within 5 seconds. Each multicast LBM uses a unique transaction ID. Transaction IDs in the same MEP must be unique within 1 minute.

After receiving a multicast LBM, each MEP on the receiver checks the multicast LBM and sends a unicast LBR. If the transaction ID of an LBR is not contained in the transaction ID list stored in the MEP that initiates loopback, the LBR is invalid and is discarded.

The MIP transparently transmits multicast LBMs. The MIP discards a received LBR that is directed to the MIP itself because the LBR is invalid.

4.3.3 ETH-LT

The Ethernet linktrace function resumes the adjacency relationship and locates faults.

By default, after sending an ETH-LT request message, a MEP receives an ETH-LT response message within 5 seconds. The MIP or MEP that receives the ETH-LT request message replies with an ETH-LT response message.

The PDU used for carrying an ETH-LT request message is an LTM, and the PDU used for carrying an ETH-LT response message is an LTR. An LTM PDU contains the following important fields:

- **Transaction ID:** indicates the transaction ID of an LTM PDU. The receiver copies the transaction ID to the LTR PDU. The source MEP checks whether the LTR corresponds to an LTM based on the transaction ID. Each LTM frame uses a unique transaction ID. Transaction IDs in the same MEP must be unique within 1 minute.
- **TTL:** The network element that receives the LTM subtracts 1 from the value of the received TTL and then copies the new value to the LTM that needs to be forwarded to the next hop. When an MIP receives an LTM whose TTL is 1, the MIP does not forward the LTM. The MIP discards any LTM whose TTL is 0.
- **OriginMAC:** indicates the MAC address of the MEP that initiates the LTM request. This field has a length of six bytes. The MIP copies this field to the LTM that needs to be forwarded to the next hop.
- **TargetMAC:** indicates the MAC address of the destination end point. This field has a length of 6 bytes. The MIP copies this field to the LTM that needs to be forwarded to the next hop.

A MEP never forwards any LTMs.

4.3.4 ETH-AIS

A MEP periodically sends Ethernet Alarm Indication Signal (ETH-AIS) frames in the opposite upstream to its peer MEPs at the specified customer MEG levels when a fault is detected. The AIS is used to notify higher layers of a low-layer fault and report an alarm. ETH-AIS is not used on Spanning Tree Protocol (STP) networks because STP has its own fault recovery mechanisms.

AIS can be sent by MEPs, including server MEPs, at the customer MEG levels when a fault is detected. For example, when abnormal signals occur during ETH-CC or the ETH-CC function is disabled, an AIS or LCK is sent downstream.

For a point-to-point ETH connection, a MEP corresponds to only a single peer MEP. Therefore, when the MEP receives an AIS, the MEP knows which peer MEP requires alarm suppression.

For multi-point Ethernet connectivity, a MEP cannot determine the faulty server-layer entity or the subsets of peer MEPs that require alarm suppression after the MEP receives an AIS. Therefore, a MEP suppresses the alarms of all peer MEPs irrespective of whether connectivity exists after the MEP receives an AIS.

A MEP performs detection immediately after receiving an AIS and suppresses the alarms relevant to all the peer MEPs. If the MEP does not receive any AIS frame within a period 3.5 times the AIS transmission interval, the MEP resets the AIS status. A MEP starts to generate LOC alarms again after the MEP detects LOC faults in case of no AISs.

The AIS transmission interval is transferred through the Flag field of an AIS PDU. Table 4-5 describes the values of the AIS/LCK transmission interval.

Table 4-5 Values of the AIS/LCK transmission interval

| Mark [3:1] | Interval | Description |
|------------|----------|---|
| 000-011 | Invalid | The value is invalid for an AIS PDU or LCK PDU. |
| 100 | 1s | One frame is transmitted per second. |
| 101 | Invalid | The value is invalid for an AIS PDU or LCK PDU. |
| 110 | 1min | One frame is transmitted per minute. |
| 111 | Invalid | The value is invalid for an AIS PDU or LCK PDU. |

In ITU-T Y.1731, the AIS transmission interval of 1 second is recommended. The first AIS frame is sent immediately after a fault is detected. The MIP transparently transmits AISs.

4.3.5 ETH-RDI

The Ethernet remote defect indication (ETH-RDI) function is performed to notify the peer MEPs of faults. The PDU used for carrying an RDI is a CCM frame, as shown in Figure 4-4. The MIP transparently transmits RDIs.

The ETH-RDI function supports the following applications:

- Single-end error management
- Remote-end performance monitoring, indicating faults that have occurred on the remote end

A faulty MEP sends an RDI. After receiving the RDI, the peer MEP knows that the peer MEP is faulty. For multi-point ETH connectivity, a MEP cannot determine the subsets of the peer MEPs that are faulty and send RDIs.

- RDIs contain an RDI transmission interval field. This field depends on the applications, and the value is set to be the same as that of the ETH-CC transmission interval.

For a point-to-point connection, a MEP can reset the RDI status after receiving the first CCM whose RDI field is deleted from the peer MEP. For multi-point ETH connectivity, a MEP can reset the RDI status only after receiving the CCM frames whose RDI fields are deleted from the entire peer MEPs in the MEP list.

4.3.6 ETH-LCK

The Ethernet lock signal (ETH-LCK) frames are sent to notify locked manageability and subsequent data service interruption. Through the ETH-LCK function, a MEP that receives the LCK can distinguish a fault from a manageability lock action of a server-layer (sublayer) MEP. One application scenario in which manageability of a MEP needs to be locked is the ETH-Test performed after service interruption. For details about the ETH-Test, see 4.3.7 ETH-Test.

When a MEP is manually locked, the MEP periodically sends LCK frames at the corresponding MEG level in the opposite upstream direction to its peer MEPs. The MEP stops sending frames only after the fault is rectified. The MIP transparently transmits LCK information.

The LCK transmission interval is the same as the AIS transmission interval. For details, see Table 4-5. If a MEP does not receive any LCK frames within a period 3.5 times the LCK transmission interval after detecting the LCK status, the MEP resets the LCK status.

An LCK PDU contains LCK transmission interval field. This field indicates the transmission interval of an LCK frame. The value is indicated by the Flag field.

4.3.7 ETH-Test

The Ethernet test signal (ETH-Test) function is used to perform unidirectional online or offline diagnosis and tests such as bandwidth flux check, frame loss check, and bit error code check.

During an offline test, data service flows of customers are interrupted in the entity to be diagnosed. The MEP that is configured for the offline test sends LCK frames at the customer MEG levels in the direction where TST frames are received.

During an online test, service data flows are not interrupted. A Test message is sent using limited bandwidth. In this case, the transmission rate of the Test message is predefined.

A TST PDU contains the following important fields:

- Sequence Number: Specified sequence number carried in each sent TST frame. Each TST frame has a unique sequence number. Sequence numbers in the same MEP must be different within 1 minute.
- Test TLV: This field is optional. The length and content can be configured in the MEP. The content can include a test code type and a random checksum. The test code can be a pseudo-random bit sequence (PRBS) that is equal to $2^{31}-1$ or an all-0 code.

During a test, a MEP inserts a Test message and sends it to the peer MEPs. The MIP transparently transmits Test messages.

4.4 OAM Performance Monitoring

The OAM performance monitoring function can be performed to measure different performance parameters. In this section, the performance parameters are used for the points

defined for point-to-point Ethernet connections. MEF 10 defines the following performance parameters:

- Frame loss rate (FLR)
- Frame delay (FD)
- Frame delay variation (FDV)
- Throughput

4.4.1 ETH-LM

The implementation of the frame loss measurement (ETH-LM) is as follows: A MEP sends an ETH-LM to the peer MEPs and receives ETH-LMs from the peer MEPs. Frame loss (FL) measurement is performed on each MEP to determine the time period during which the bidirectional service is unavailable. If any of the two directions is unavailable, the bidirectional service is unavailable.

The FLR is defined as the ratio of the number of non-transferred service frames to the total number of service frames within in the interval T, in percentage. The number of non-transferred service frames refers to the difference between the number of service frames that reach the ingress ETH connection point and the number of service frames that are transferred to the egress ETH connection point in case of point-to-point ETH connection.

An ETH-LM contains the ETH-LM transmission interval. The default transmission interval is 100 ms.

The MIP transparently transmits ETH-LM frames.

The ETH-LM can be performed in the following modes:

- Dual-end ETH-LM
The dual-end ETH-LM mode is adopted for error management. Each MEP terminates ETH-LM messages. The local and remote FL measurements are performed on each MEP. The dual-end ETH-LM is performed to monitor the performance at the same priority with the ETH-CC.
An ETH-LM message is carried in a CCM PDU. The CCM PDU transmission interval is the same as the CCM transmission interval that is set by sending MEP for performance monitoring.
- Single-end ETH-LM
In this mode, a MEP sends an ETH-LM request message to the peer MEPs and receives ETH-LM response messages from the peer MEPs to perform FL measurement.
A single-end ETH-LM request message is carried in an LMM frame, and a single-end ETH-LM response message is carried in an LMR frame. The LMM transmission interval is different from that of a dual-end ETH-LM. The LMMs are sent based on the value of the TxFCf counter.

4.4.2 ETH-DM

The frame delay measurement (ETH-DM) function is used for on-demand OAM. The ETH-DM function is performed to measure the FD and FDV. The FD refers to the interval between the time when the first bit of a frame is transmitted by a source node and the time when the last bit of the frame is received by the same source node after the frame is looped back. Frame loopback is performed by the destination node of the test.

In this method, the loopback function is performed for measuring the delay of the loop formed by each request frame and the corresponding response frames or the bidirectional frame delay.

A sender sends an OAM request message with a timestamp. A recipient copies the timestamp in the sent OAM request message. You can calculate the difference between the time when an OAM response message is received and the original timestamp carried in the OAM response message. In this way, the loop delay is calculated.

The FD and FDV measurements can be performed on each MEP. The MIP transparently transmits ETH-OM frames.

The ETH-DM can be performed in two modes: unidirectional ETH-DM and bidirectional ETH-DM.

4.4.3 FDV

The FDV function is performed to measure the delay variation between the frames of instances of the same class of service (CoS) level on a point-to-point ETH connection. In this method, the delay of the loop formed by each request frame and the corresponding response frames or the bidirectional frame delay is measured. Within the measurement period, the sender of a request message records the maximum delay (FDmax) and the minimum delay (FDmin).

The FDV is calculated as follows:

$$FDV = FD_{max} - FD_{min}$$

The information elements of an FDV in OAM data include the sequence number and the request timestamp.

4.4.4 Throughput Measurement

Data frames are sent at incremental rates. The maximum rate can be reached. The frame receiving percentages are recorded in curves. The rate at which frames become discarded is reported. In normal cases, this rate depends on the data frame length.

In IEEE 802.1ag, the unicast ETH-LB messages such as LBM or LBR frames that carry data fields and the ETH-Test messages such as TST frames that carry data fields can be used for the throughput measurement. A MEP can insert TST or LBM frames with the specified length and code type at a certain rate to perform unidirectional or bidirectional flux measurement.

4.4.5 Usability

Usability indicates the availability of an ME, expressed as ratio. Usability is defined as the ratio of the time period during which an ME is available to the total service time period. The time period during which an ME is available is the time when the service is within the boundary of the FL, FD, and FDV. The time period during which an ME is unavailable indicates the time when the service exceeds any of the thresholds of the FL, FD, and FDV. The thresholds are defined by the CoS. The measurement is performed on the basis of the FL, FD, and FDV. The time period during which an ME is available (for example, 24 hours) can be divided into multiple measurement intervals, for example, one minute. The FL, FD, and FDV are measured within each measurement interval. If the threshold for a certain service type is exceeded, the ME is considered unavailable in this measurement interval; otherwise, the ME is available.

Usability is calculated as follows:

$$\text{Usability} = \text{Number of measurement intervals when the ME is available} / \text{Number of total measurement intervals} \times 100\%$$

4.5 Comparison Between ITU-T Y.1731 and IEEE 802.1ag

Table 4-6 describes the comparison between ITU-T Y.1731 and IEEE 802.1ag.

Table 4-6 Comparison between ITU-T Y.1731 and IEEE 802.1ag

| OAM Function | Y.1731 | 802.1ag | Description |
|----------------------------|---------|-----------|--|
| Connectivity check (CC) | Defined | Defined | The minimum and maximum CCM intervals of ITU-T Y.1731 are different from those of IEEE 802.1ag. |
| Unicast loopback | Defined | Defined | |
| Multicast loopback | Defined | Undefined | |
| Multicast linktrace | Defined | Defined | |
| ETH-AIS | Defined | Undefined | Ethernet alarm indication |
| ETH-RDI | Defined | Defined | |
| ETH-LCK | Defined | Undefined | LCK frames are sent to indicate the locked status. Data transmission is interrupted. |
| ETH-TST | Defined | Undefined | TST frames are sent to perform throughput measurement and mis-sequencing frame measurement. |
| ETH-APS | Defined | Undefined | ITU-T G.8031/Y.1342 is referenced. |
| ETH-USM/USR | Defined | Undefined | Channel maintenance function. Remote management of MEPs is supported. |
| ETH performance monitoring | Defined | Undefined | This function is performed to measure performance parameters such as the FLR, FD, FDV, and throughput. |
| ETH-EXM/EXR | Defined | Undefined | This function is used in experiments and can be customized. |
| ETH-VSM/VSR | Defined | Undefined | This function is a proprietary OAM function and can be customized by users. |

As described in Table 4-6, ITU-T Y.1731 focuses on aspects related to the AIS, APS, and performance manager. The following functions are implemented in both ITU-T Y.1731 and IEEE 802.1ag:

1. Connectivity check and alarm indication (CCMs or CCMs with RDIs)
2. Linktrace (MAC tracert and LT packets)
3. Link discovery (MAC ping and LB packets)

The protocol packets such as CCMs, LB packets, and LT packets of the ITU-T Y.1731 are similar to but are not exactly the same as those of the IEEE 802.1ag. The protocol packets of ITU-T Y.1731 and IEEE 802.1ag are incompatible.

5 Fault Association of Ethernet OAM

5.1 Association Between EFM OAM and an Interface

When an interface enabled with EFM OAM detects a connectivity fault between the interface and the remote interface, the OAM management module shuts down the interface for 7 seconds and then restarts the interface.

5.2 Association Between EFM OAM and BFD

When the EFM OAM module detects a fault, the OAM management module sends the fault message to BFD at the other side through the interface. When BFD detects a fault, BFD sends the fault message to EFM OAM through the interface.

- EFM OAM sends fault messages to BFD.
- BFD sends fault messages to EFM OAM.
- EFM OAM and BFD perform bidirectional transmission of fault messages.

5.3 Association Between EFM OAM and MPLS OAM

When the EFM OAM module detects a fault, the OAM management module sends the fault message to MPLS OAM module through the interface. When MPLS OAM detects a fault, the OAM management module sends the fault message to EFM OAM module through the interface.

- EFM OAM sends fault messages to MPLS OAM.
- MPLS OAM sends fault messages to EFM OAM.

5.4 Association Between Ethernet CFM and EFM OAM

When the Ethernet CFM module detects a fault in an MA, the OAM management module sends the fault message to the peer device enabled with EFM OAM through the interface. When the EFM OAM module detects a fault, the OAM management module sends the fault message to the MA through the interface.

- Ethernet CFM sends fault messages to EFM OAM.
- EFM OAM sends fault messages to Ethernet CFM.
- Ethernet CFM and EFM OAM perform bidirectional transmission of fault messages.

5.5 Association Between Ethernet CFM and an Interface

When a MEP detects a connectivity fault between the MEP and a specified RMEP within the same MA, the OAM management module performs shuts down the interface on which the MEP resides for 7 seconds and then restarts the interface.

5.6 Association Between Ethernet CFM and Ethernet CFM

When the Ethernet CFM module detects a fault in an MA, the OAM management module sends the fault message to the MA at the other side through the binding relationship.

- Ethernet CFM at one side sends fault messages to Ethernet CFM at the other side.
- Ethernet CFMs at both sides perform bidirectional transmission of fault messages.

5.7 Association Between Ethernet CFM and BFD

When the Ethernet CFM module detects a fault in an MA, the OAM management module sends the fault message to BFD at the other side through the binding. When BFD detects a fault, BFD sends the fault message to the MA through the binding relationship.

- Ethernet CFM sends fault messages to BFD.
- BFD sends fault messages to Ethernet CFM.
- Ethernet CFM and BFD perform bidirectional transmission of fault messages.

5.8 Association Between Ethernet CFM and MPLS OAM

When the Ethernet CFM module detects a fault in an MA, the OAM management module sends the fault message to MPLS OAM through the binding relationship. When MPLS OAM detects a fault, the OAM management module sends the fault message to the MA through the binding relationship.

- Ethernet CFM sends fault messages to MPLS OAM.
- MPLS OAM sends fault messages to Ethernet CFM.

5.9 Association Between Ethernet CFM and Smart Link

When the Ethernet CFM module detects a fault in an MA, the OAM management module sends the fault message to Smart Link at the other side through the binding. Then the active/standby switchover of Smart Link is triggered.

6 Ethernet OAM Application Scenarios

This chapter describes the Ethernet OAM application scenarios of switches.

6.1 Combination of EFM, CFM, and Smart Link

6.1.1 Introduction

Enterprises connect to aggregation switches through DSLAMs. To improve reliability of the network, each aggregation switch is dual homed to core switches through two uplinks in a Smart Link group.

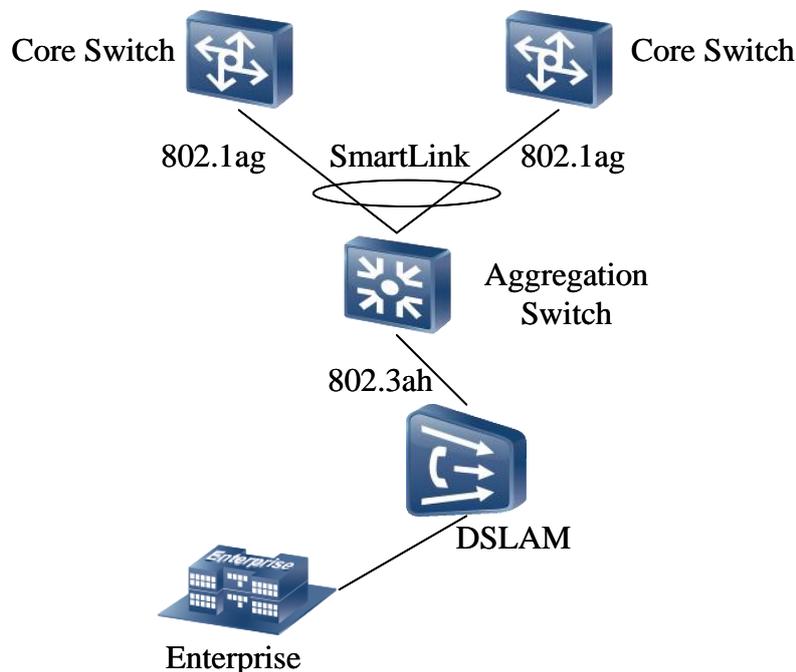
IEEE 802.1ag link detection can be enabled on the active and standby links of the Smart Link group. In this way, the link status can be detected quickly. The hardware engines of the switches are used to implement fault detection within 10 ms so that an active/standby switchover of the Smart Link group can be triggered.

In addition, IEEE 802.3ah can be enabled between the aggregation switch and the DSLAM. In this way, the link status can be monitored, and a diagnosis measure such as a loopback test can be performed.

Through the enabling of the fault notification functions of the IEEE 802.1ag and IEEE 802.3ah, remote ends can be notified of any occurring faults.

6.1.2 Networking

Figure 6-1 Application scenario I



6.2 Combination of EFM/CFM, MPLS OAM, and BFD

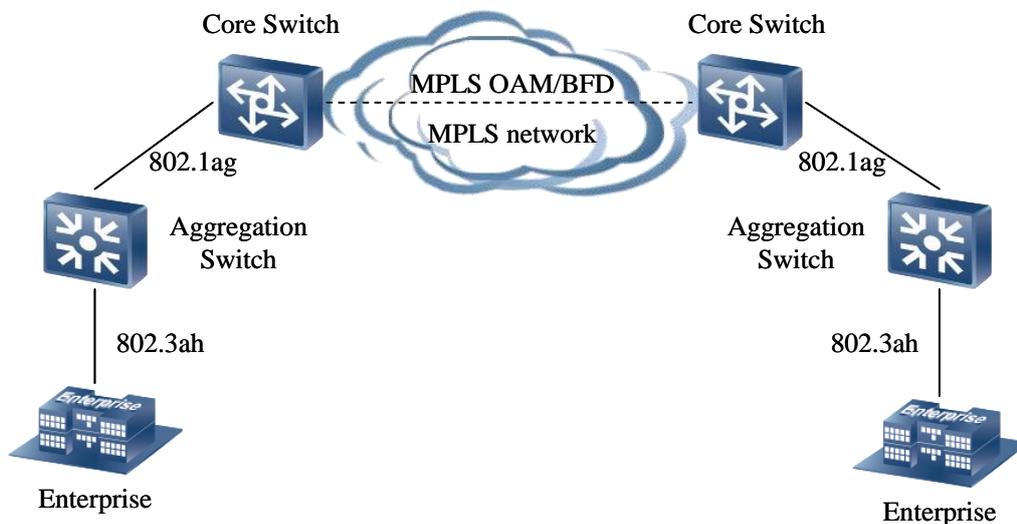
6.2.1 Introduction

User services are connected to the network through aggregation switches. After being transmitted to core switches, service packets are forwarded to other service access points through the MPLS network. IEEE 802.3ah detection is enabled between aggregation switches and users. IEEE 802.1ag detection is enabled between aggregation switches and core switches. MPLS OAM or BFD is enabled on the MPLS network for link detection. In this way, a solution to end-to-end multi-level link detection and protection is provided.

The association of the IEEE 802.1ag detection with MPLS OAM or the BFD detection can be enabled to implement the association detection of the services that are connected to the core network.

6.2.2 Networking

Figure 6-2 Application scenario II



IEEE 802.1ag is enabled between aggregation switches and core switches for link detection. MPLS OAM or BFD is enabled between core switches for link detection. When a service link fault occurs between an aggregation switch and a core switch, the ETH OAM module notifies MPLS OAM or the BFD module of the lower-layer link fault.

A

Abbreviation

| Abbreviation | Full Name |
|---------------------|--|
| AIS | Alarm Indication Signal |
| CCM | Continuity Check Message |
| DSAP | Domain Service Access Point |
| EVC | Ethernet Virtual Connection |
| FD | Frame Delay |
| FDV | Frame Delay Variety |
| FL | Frame Loss |
| LBM | Loopback Message |
| LBR | Loopback Response |
| LTM | LinkTrace route Message |
| LTR | LinkTrace route Response |
| MA | Maintenance Association |
| MAID | Maintenance Association Identifier |
| MD | Maintenance Domain |
| ME | Maintenance Entity |
| MEG | Maintenance Entity Group |
| MEP | Maintenance association End Point |
| MIP | Maintenance association Intermediate Point |
| ISAP | Intermediate Service Access Point |
| OAM | Operations, Administration and Maintenance |
| RDI | Remote Defect Indication |