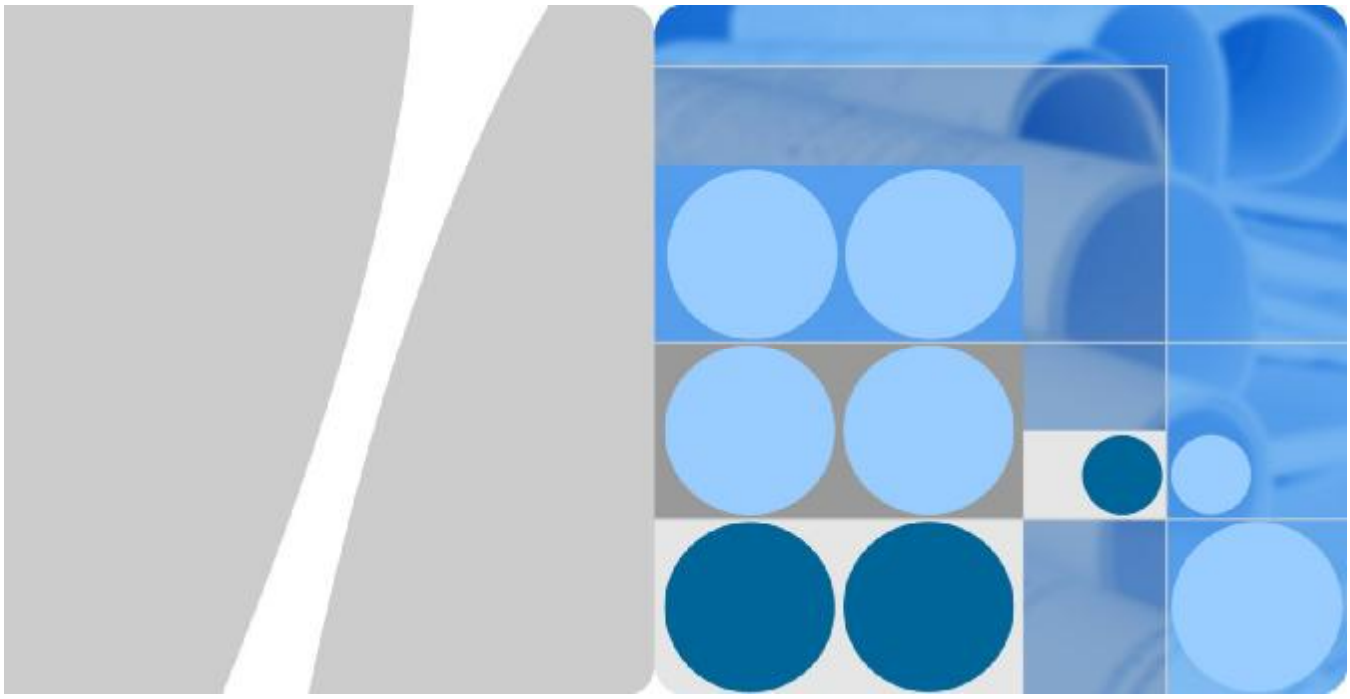

S Series Switch

Ethernet OAM Technology White Paper

Issue 3.0

Date 2015-09-15

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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://e.huawei.com

Contents

Technical White Paper for Etherent OAM

Abstract: This document describes the background, technical points, and application scenarios of Ethernet OAM.

Keywords: OAM, 802.1ag, 802.3ah, Y.1731, EFM, CFM

1 Introduction to Ethernet OAM

1.1 Overview

Easy-to-use Ethernet techniques support good bandwidth expansibility on low-cost hardware. With these advantages, the Ethernet has been widely used and becomes an important technology for communications networks. Compared with other networks, the Ethernet is characterized by the advantage in low construction cost. With continuous development of the Ethernet, especially traffic transmission from the local area network (LAN) to the wide area network (WAN), the network management and maintenance become more and more important. Traditional Ethernet can hardly provide end-to-end (E2E) service management, fault detection, and performance monitoring, so it is difficult to locate and rectify faults occurring on the traditional Ethernet. The maintainability and operability are low. To address the preceding issues, the International Telecommunication Union (ITU), Institute of Electrical and Electronics Engineers (IEEE), and Metro Ethernet Forum (MEF) had defined protocols and standards for Ethernet operations, administration, and maintenance (OAM), including IEEE 802.1ag and 802.1ah, ITU-T Y.1731, and Ethernet Local Management Interface (E-LMI) of MEF. These protocols and standards complement each other to provide E2E service operation management and maintenance capabilities. Figure 1-1 shows the architecture of Ethernet OAM.

Figure 1-1 Architecture of Ethernet OAM

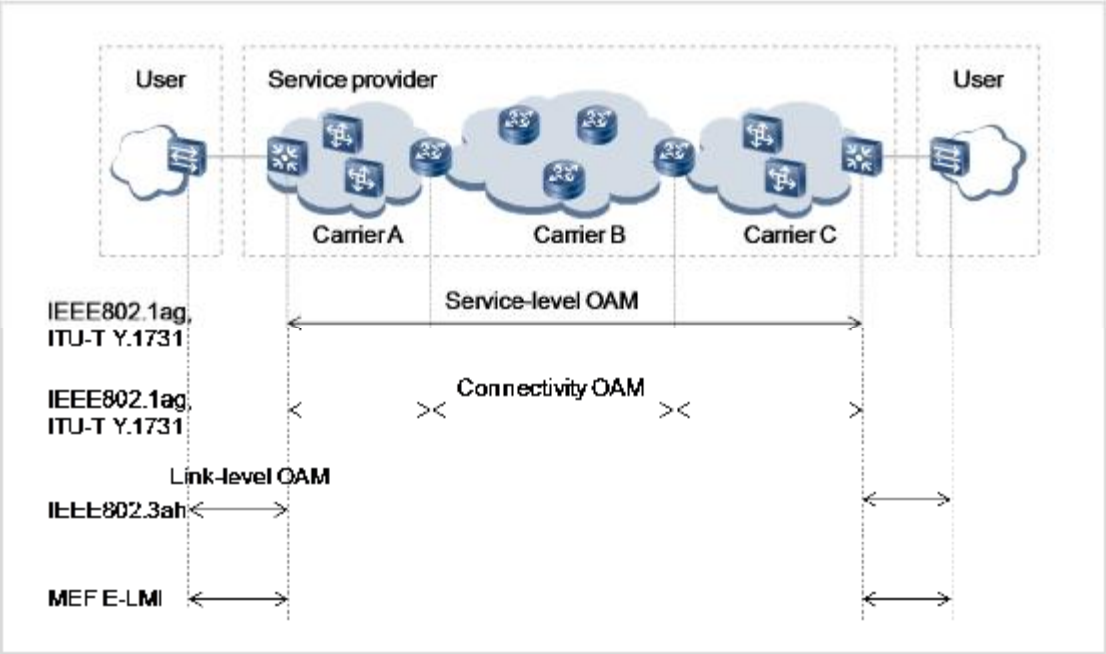

Ethernet OAM provides the following functions:

- Fault management
  - Detect the network connectivity by periodically or manually sending detection messages.
  - Acknowledge and locate faults on an Ethernet network by using methods similar to the ping and traceroute on an IP network.
  - Work with protection switching protocols to trigger protection switching when connectivity faults are detected. This implements service interruption within 50 ms to achieve carrier-class reliability.
- Performance management

  Ethernet OAM measures network transmission parameters including the packet loss ratio, delay, and jitter, and collects statistics on various types of packets. Performance management is often implemented on access devices. With performance management tools, carriers can monitor the network running status and locate faults through the network management system (NMS). The carriers can then check whether forwarding capabilities of networks comply with the Service Level Agreement (SLA) signed with users.

Huawei switches provide IEEE 802.3ah OAM and IEEE 802.1ag OAM. IEEE 802.3ah OAM and IEEE 802.1ag OAM calculate the performance indicators such as the delay based on the performance evaluation methods defined in the ITU-T Y.1731 protocol. Huawei switches enabled with Ethernet OAM provide 3.3-ms fault detection, which is the fastest in industry. The major standards are described as follows.

1.2 Overview of Ethernet OAM

IEEE 802.3ah, also known as Ethernet in the First Mile (EFM), is a link-level OAM mechanism and defines the specifications of the Ethernet physical layer and OAM used for user access. EFM provides link connectivity check, link fault monitoring, remote fault notification, and remote loopback for a link between two directly connected devices.

IEEE 802.1ag, also known as Connectivity Fault Management (CFM), is a network-level OAM mechanism that provides Ethernet OAM functions, such as continuity check (CC), loopback (LB), and linktrace (LT). CFM applies to large-scale, end-to-end networks.

ITU-T Y.1731 is a service-level OAM mechanism that provides fault management functions similar to IEEE 802.1ag. In addition to E2E connectivity check, LB, LT, Y.1731 provides diagnosis testing and performance management functions such as frame loss measurement, frame delay measurement, frame jitter measurement, and throughput measurement.

2 EFM Implementation

2.1 Basic Concepts

2.1.1 Protocol Packet

EFM works at the data link layer and uses protocol packets called OAM Protocol Data Units (OAMPDUs). After interfaces of two directly connected devices are enabled with EFM, the devices periodically exchange OAMPDUs to report the link status so that network administrators can effectively manage networks. The interface enabled with EFM is called an EFM entity.

EFM is a slow protocol defined by IEEE 802.3. For a slow protocol, a maximum of 10 frames of protocol packets are transmitted per second, the destination multicast MAC address is 0x0180-C200-0002, and the Ethernet type is 0x8809. The packets does not carry VLAN IDs and cannot be directly forwarded by network bridges. OAMPDUs can be transparently transmitted by transmission devices or devices enabled with Layer 2 protocol transparent transmission to detect the transmission link between two EFM entities. Figure 2-1 shows the OAMPDU format and common types of OAMPDUs. Table 2-1 describes fields in an OAMPDU.

Figure 2-1 OAMPDU format

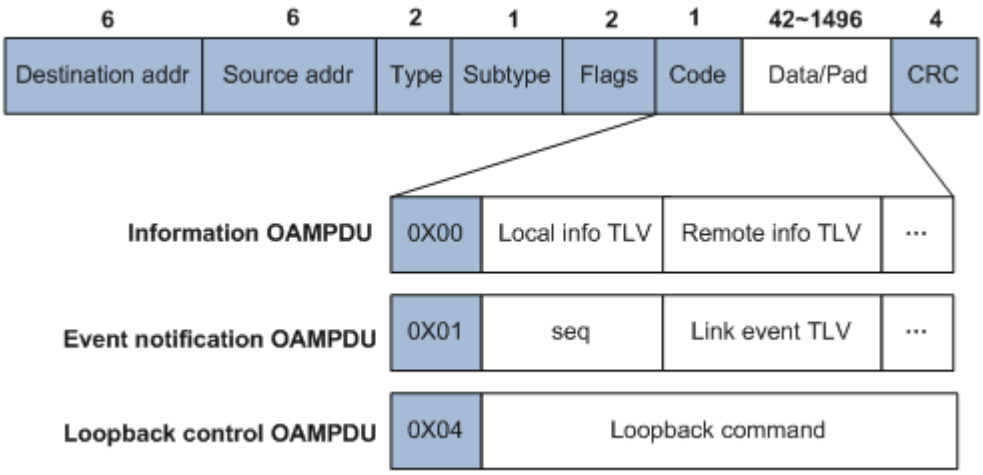

Table 2-1 Fields in an OAMPDU

Field	Description
Destination addr	The destination MAC address is 0x0180-C200-0002. Slow protocol packets cannot be forwarded by network bridges or over multiple devices.
Source addr	Source address, which is a unicast MAC address of an interface on the transmit end. If no interface's MAC address is specified on the transmit end, the bridge MAC address of the transmit end is used.
Type	Slow protocol type, which has a fixed value of 0x8809.
Subtype	Subtype of a slow protocol. The value is 0x03, indicating that the slow subprotocol is EFM.
Flags	Status of an EFM entity: - Remote Stable - Remote Evaluating - Local Stable - Local Evaluating - Critical Event - Dying Gasp - Link Fault
Code	Message code value, which specifies a specific type of OAMPDU. Table 2-2 lists types of OAMPDUs.

Table 2-2 OAMPDU types

Code Value	OAMPDU Type	Function
0x00	Information OAMPDU	Used by local and remote EFM entities to exchange the Flags field, check connectivity periodically, and advertise fault information.
0x01	Event notification OAMPDU	Used to monitor links. If an errored frame event or errored code period event occurs on an interface, the interface sends an Event Notification OAMPDU to notify the remote interface of the event.
0x04	Loopback control OAMPDU	Used to enable or disable the remote loopback function.

2.1.2 Connection Modes

EFM supports two connection modes: active and passive. An EFM connection can only be initiated by an OAM entity working in active mode. An OAM entity working in passive mode waits for receiving a connection request from its remote entity. Table 2-3 lists capabilities for processing OAMPDUs in the two modes.

Table 2-3 Capabilities for processing OAMPDUs in active and passive modes

Capability	Active Mode	Passive Mode
Initiate a connection request by sending an Information OAMPDU during the discovery process.	Supported	Not supported
Respond to a connection request during the discovery process.	Supported	Supported
Send Information OAMPDUs.	Supported	Supported
Send Event Notification OAMPDUs.	Supported	Supported
Send Loopback Control OAMPDUs.	Supported	Not supported
Respond to Loopback Control OAMPDUs.	Supported (The remote EFM entity must work in active mode.)	Supported

2.1.3 Link Event

Link events defined in EFM fall into minor and critical link events. Minor link events are used to monitor link performance, as described in Table 2-4:

- Errored Symbol Period Event: If the number of symbol errors that occur on a device interface in a given period of time reaches or exceeds a configured threshold, the device generates an Errored Symbol Period Event.
- Errored Frame Event: If the number of frame errors that occur on a device interface in a given period of time reaches or exceeds a configured threshold, the device generates an Errored Frame Event.
- Errored Frame Seconds Summary Event: An errored frame second is a 1-second interval during which at least one errored frame is detected. If the number of errored frame seconds that occur on a device interface in a given period of time reaches or exceeds a configured threshold, the device generates an Errored Frame Second Summary Event. When one or more errored frames are detected within 1 second, the switch calculates one errored frame second. The number of errored frame seconds in a given period of time is different from the number of frame errors. The Errored Frame Event involves the number of detected error frames within a given period of time. For example, the detection time of errored frame seconds is 60s and the threshold is 20. When 20 errored frame seconds are detected within 60s, an Errored Frame Seconds Summary Event is triggered. If error frames are generated every second, a maximum of 60 errored frame seconds are detected in 60s.

Table 2-4 Minor link events

Minor Link Event	Description	Usage Scenario
Errored Symbol Period Event	The switch generates an Errored Symbol Period Event, advertises the event to the remote device, and sends a trap to the NMS.	This event helps the switch to detect code errors during data transmission at the physical layer.

Minor Link Event	Description	Usage Scenario
Errored Frame Event	The switch generates an Errored Frame Event, advertises the event to the remote device, and sends a trap to the NMS.	This event helps the switch to detect error frames that occur during data transmission at the data link layer.
Errored Frame Seconds Summary Event	The switch generates an Errored Frame Seconds Summary Event, advertises the event to the remote device, and sends a trap to the NMS.	This event helps the switch to detect errored frame seconds that occur during data transmission at the data link layer. The statistics collection methods of the Errored Frame Event and Errored Frame Seconds Summary Event are different.

Critical link events are used for remote fault detection. Table 2-5 describes the critical link events.

Table 2-5 Critical link events

Critical Link Event	Description
Link fault	If a loss of signal (LoS) error of the remote EFM entity occurs because a physical link fails, the local device sends a trap to the NMS.
Dying gasp	If an unexpected status change or event occurs, for example, power-off of the remote device and restart of the local device or card, the local device sends a trap to the NMS.
Critical event	If an unidentified critical event occurs because a fault is detected using association between the remote EFM entity and a specific feature, the local device sends a trap to the NMS. Remote EFM entities can be associated with protocols, including Bidirectional Forwarding Detection (BFD), Connectivity Fault Management (CFM), and Multiprotocol Label Switching (MPLS) OAM.
Link loss	If a loss of signal (LoS) error of the remote EFM entity occurs because the interval at which OAMPDUs are sent elapses, the local device sends a trap to the NMS.

2.2 Working Mechanism

EFM supports the following functions: OAM discovery, link monitoring, fault notification, and remote loopback. The following example illustrates EFM implementation on the network shown in Figure 2-2. The customer edge (CE) is an edge device on the customer side and PE1 is a network-side device. EFM is deployed between the CE and PE1 to remotely monitor link connectivity and quality.

Figure 2-2 Typical EFM networking

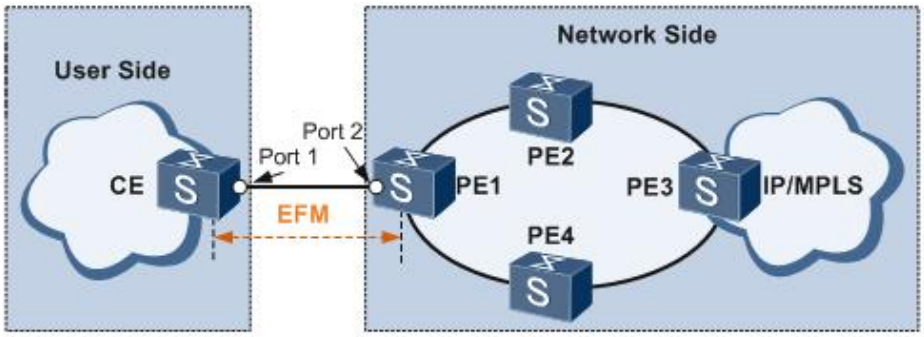

2.2.1 EFM Discovery

During the discovery process, a local EFM entity discovers and establishes a stable EFM connection with a remote EFM entity. Figure 2-3 shows the discovery process.

Figure 2-3 EFM discovery

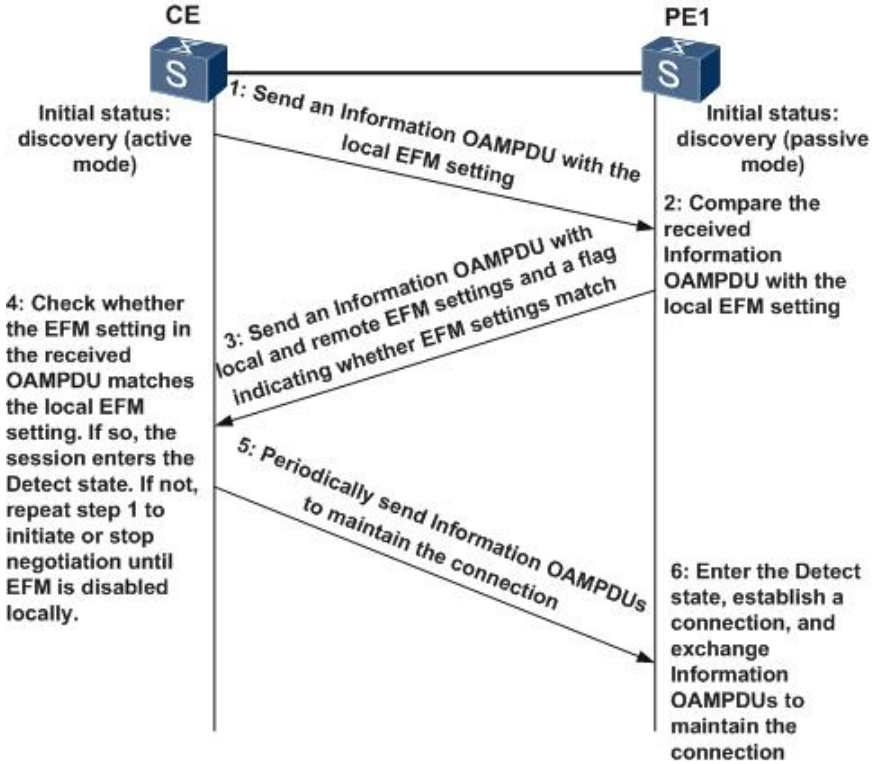

After the EFM connection is established, EFM entities at both ends periodically exchange Information OAMPDUs to monitor link connectivity. The interval at which Information OAMPDUs are sent is also known as an interval between handshakes. If an EFM entity does not receive Information OAMPDUs from the remote EFM entity within the timeout interval, the EFM entity considers the connection interrupted and sends a trap to the NMS. Generally, the interval between handshakes is 1s and the timeout interval is 5s. EFM connection setup is used to automatically monitor physical link connectivity. EFM is a slow protocol and provides second-level detection and timeout, so it does not need to provide 3.3-ms fast fault detection.

2.2.2 Link Monitoring

When traffic is being transmitted over physical links and network performance deteriorates, for example, the link quality is lowered, external factors affect the link, and the link is congested, it is difficult to detect faults on the Ethernet link. To resolve this problem, configure the EFM link monitoring function that detects data link layer faults in various environments. EFM entities enabled with link monitoring exchange Event Notification OAMPDUs to monitor links.

If an EFM entity receives a link event listed in Table 2-4, it sends an Event Notification OAMPDU to notify the remote EFM entity of the event and also sends a trap to the NMS. After receiving the trap on the NMS, an administrator can determine the network status and take remedial measures as needed.

2.2.3 Fault Notification

After the OAM discovery process finishes, two EFM entities at both ends connection exchange Information OAMPDUs to monitor link connectivity. The sending interval is 1s and the timeout interval is 5s by default. When traffic is interrupted because the remote EFM entity fails or becomes unavailable, the faulty EFM entity will send an Information OAMPDU carrying a critical link event listed in Table 2-5 to the local EFM entity. After receiving the notification, the local EFM entity sends a trap to the NMS. An administrator can view the trap on the NMS to determine the link status and take measures to rectify the fault. Common faults include device restart, interface card reset, physical link fault, protocol packet timeout, and faults transmitted by the OAM module (for example, fault transmission through association between EFM and BFD or CFM).

2.2.4 Remote Loopback

When a local interface sends non-OAMPDUs to a remote interface, the remote interface loops the non-OAMPDUs back to the local interface, but not to the destination addresses of the non-OAMPDUs. This is remote loopback. An EFM connection must be established to implement remote loopback.

After remote loopback is enabled, the device discards all the non-OAMPDUs, causing service interruption. It is recommended that you enable remote loopback to check link connectivity and quality before a new network is built to provide services or a link fault is rectified. The results help an administrator to take measures to minimize the impact of remote loopback on services.

The local device computes communication quality parameters such as the packet loss ratio on the current link based on the numbers of sent and received test packets. Figure 2-4 shows the implementation of remote loopback. If the local device attempts to stop remote loopback, it sends a message to instruct the remote device to disable remote loopback. After receiving the message, the remote device disables remote loopback.

Figure 2-4 Implementation of remote loopback

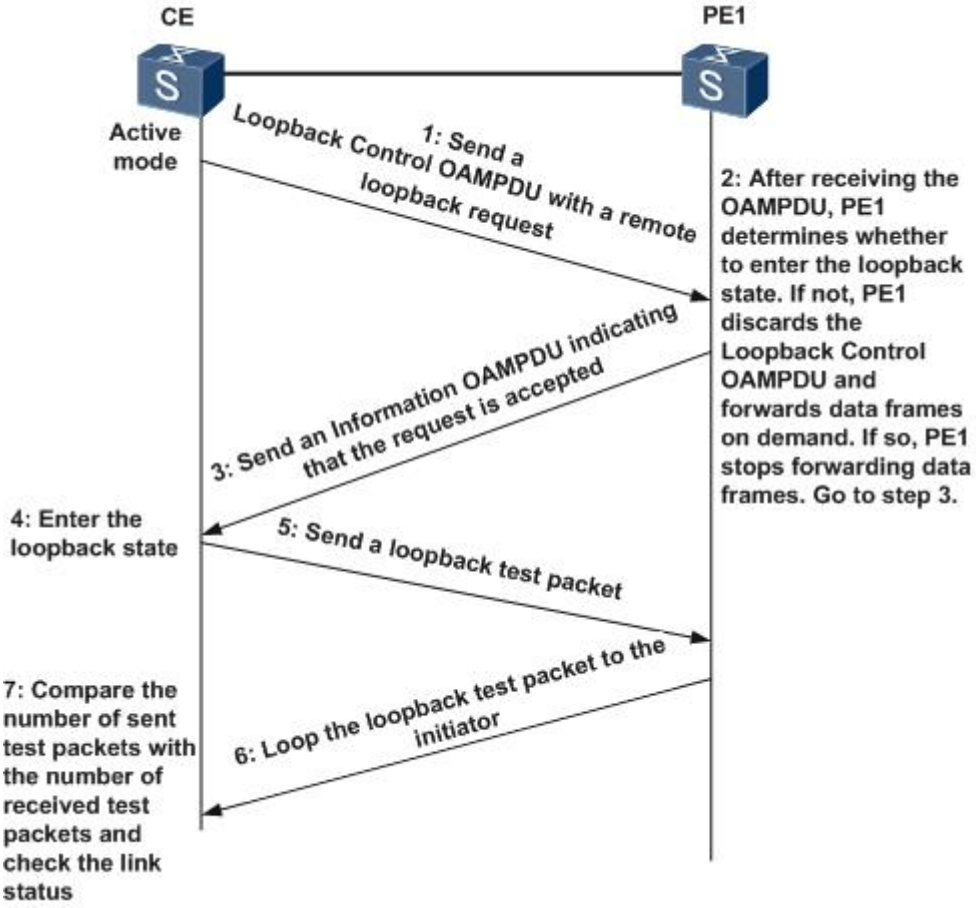

3 CFM Implementation

3.1 Basic Concepts

3.1.1 MD

A maintenance domain (MD) is a network or a part of a network where CFM is performed. Generally, an MD is managed by an Internet service provider (ISP).

To help locate faults, MDs are classified into levels 0 through 7. A larger value indicates a higher level and a larger area that the MD covers. CCMs in high-level MDs can pass through low-level MDs, but CCMS in low-level MDs cannot pass through high-level MDs. One MD can be tangential to or nested into another MD, but MDs cannot be crossed. In addition, only a lower-level MD can be nested in a higher-level MD.

As shown in Figure 3-1, MD2 is nested within MD1. CCMs in MD1 need to pass through MD2. The levels of MD1 and MD2 are set to 6 and 3, respectively. The level setting allows CCMs to pass through MD2 to monitor connectivity of MD1, but prevents CCMs from passing through MD1. Properly planned MDs help a network administrator locate faults.

Figure 3-1 MD networking

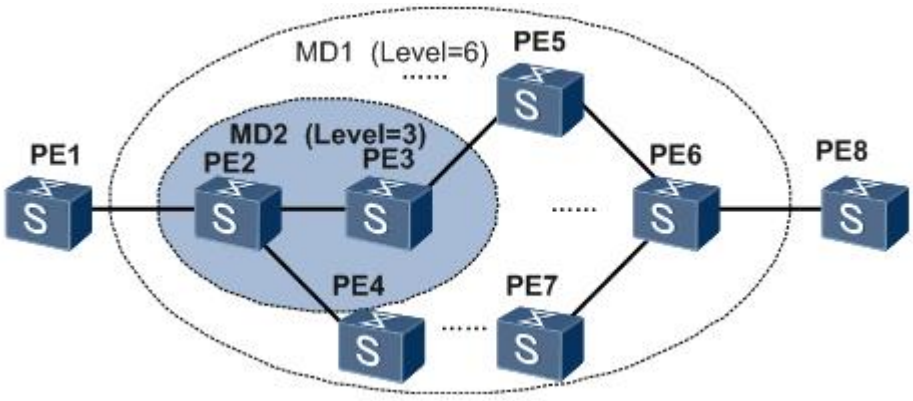

3.1.2 MA

A maintenance association (MA) is a part of an MD. One or more MAs can be configured in an MD. An MA serves a specific service such as VLAN, that is, each MA corresponds to one VLAN. CFM detects connectivity faults on each MA.

3.1.3 MEP

A maintenance association end point (MEP) determines the scope and edge of an MD and is the edge node of an MA.

For any network device running CFM, its MEP is called the local MEP. For the other devices in the same MA, their MEPs are called remote maintenance association end points (RMEPs).

3.1.4 MIP

A maintenance association intermediate point (MIP) is located on a link between two MEPs within an MD, facilitating management. More MIPs make network management and control easier.

3.1.5 CFM Protocol Packets

CFM sends different types of protocol packets to detect and locate link faults. Figure 3-2 shows common protocol packets including the continuity check message (CCM), loopback reply (LBR) message, linktrace reply (LTR) message, and linktrace message (LTM). Table 3-1 describes the functions of CFM protocol packets.

Figure 3-2 CFM protocol packet format

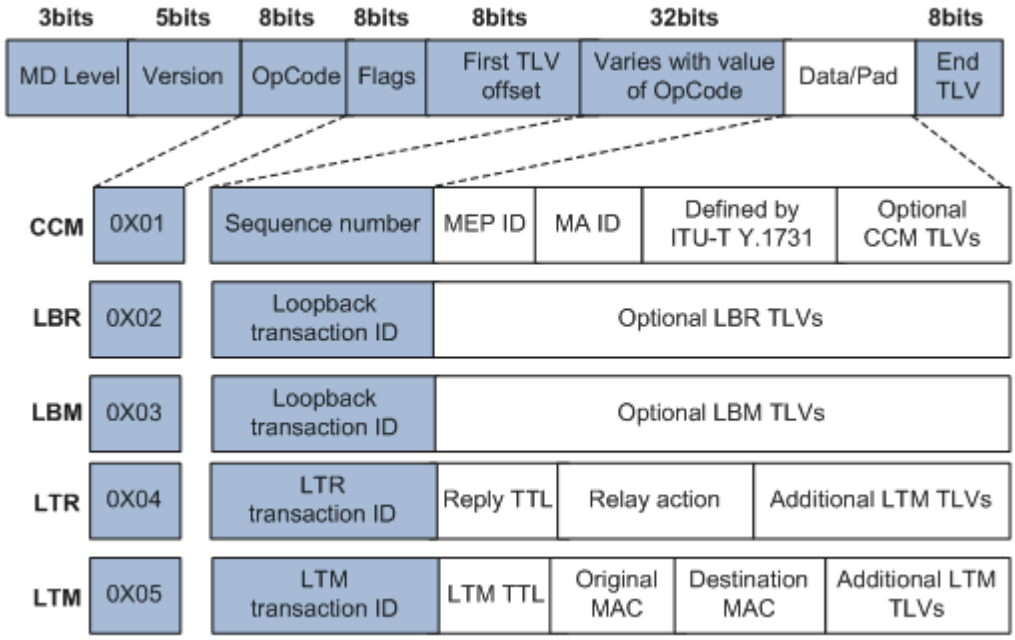

Table 3-1 Types of CFM protocol packets

OpCode Value	Packet Type	Function
0x01	CCM	Monitors end-to-end link connectivity.
0x02	LBR	Is used by the remote device to respond to an LBM sent by the local device.
0x03	LBM	Is sent by an interface that initiates loopback detection.
0x04	LTR	Is used by the remote device to respond to an LTM sent by the local device.
0x05	LTM	Is sent by an interface to initiate a linktrace test.

3.2 Working Mechanism

CFM supports continuity check (CC), loopback (LB), and linktrace (LT) functions.

3.2.1 Connectivity Check

CC monitors connectivity of links between MEPs. A MEP periodically multicasts CCMs to an RMEP in the same MA. If an RMEP does not receive a CCM within a period three times the timeout interval at which CCMs are sent, the RMEP considers the path between itself and the MEP faulty. A log is generated to provide information for fault diagnosis. You can implement loopback or linktrace to locate the fault. If CM has been associated with another protocol or an interface, services can be switched to the standby link to ensure nonstop service transmission and fast convergence.

CCMs are generated and terminated by MEPs. A MEP forwards received CCMs of a higher level but drops CCMs of a lower level or of the same level. In this manner, CCMs from a low-level MD are not spread to the MD of a higher level.

The interval at which CCMs are sent is configurable. IEEE 802.1ag defines the following intervals: 3.3 milliseconds, 10 milliseconds, 100 milliseconds, 1 second, 10 seconds, 1 minute, and 10 minutes. Huawei switches provide the following intervals: 3.3 milliseconds, 10 milliseconds, 100 milliseconds, 1 second, and 10 seconds.

Figure 3-3 CFM system

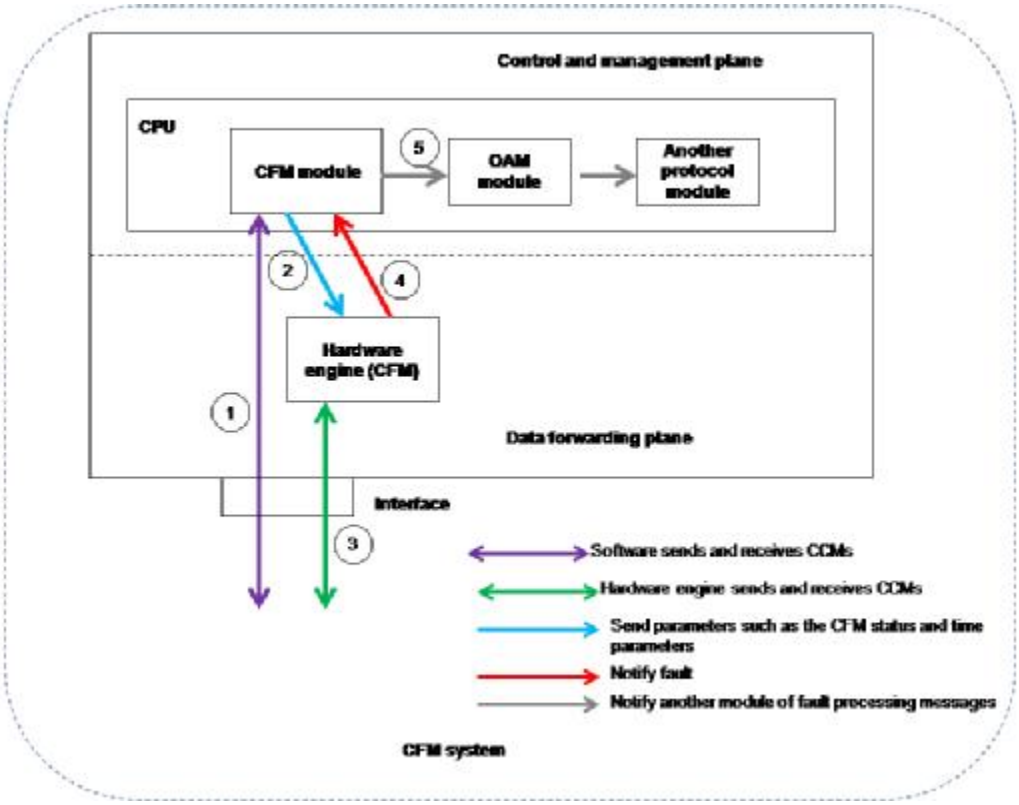

To provide fast and stable link fault detection, Huawei switches send CCMs at the minimum interval of 3.3 ms through the hardware engine. As shown in Figure 3-3, the implementation of CFM based on the hardware engine is as follows:

1. In initial state, the CFM module sends CCMs to the remote end for negotiation. When both ends of the CFM session acknowledge the configured parameters, CFM enters the Up state.
2. The CFM module sends parameters to the hardware engine. The parameters include the Up state, CCM sending interval, and timeout interval.
3. The hardware engine saves delivered entries and sets the entry status to Up. Then the hardware engine periodically sends CCMs to the outbound interface. If the CCM sending interval is 3.3 ms, the hardware engine sends CCMs at an interval of 3.3 ms and detects whether CCMs are received from the remote end within the timeout interval.
4. If the hardware engine does not receive CCMs from the remote end within three intervals, the hardware engine considers the link faulty and immediately notifies the CFM module of handling the fault.
5. The CFM module reports the log about the fault to the NMS and notifies the OAM module of the fault. If CFM has been associated with another protocol or an interface, the CFM module notifies the associated protocol or interface (CFM can be associated with EFM, BFD, MPLS OAM, ERPS, and SEP) of the fault for protection switching. Then the CFM session enters the initial state. When the link is restored, a new round of negotiation is initiated to the remote end.

3.2.2 Loopback

Loopback is also called 802.1ag MAC ping. Similar to IP ping, loopback monitors connectivity of a path between local and remote devices. IP ping can detect Layer 3 connectivity, whereas 802.1ag MAC Ping can only detect Layer 2 connectivity in a VLAN.

A MEP initiates an 802.1ag MAC ping test to monitor reachability of a MEP or MIP. The initiator and destination node must have the same level. After receiving an LBM from the MEP, the RMEP or MIP replies an LBR. Loopback helps locate a faulty node because a faulty node cannot send an LBR in response to an LBM. LBMs and LBRs are unicast packets.

Figure 3-4 Loopback networking

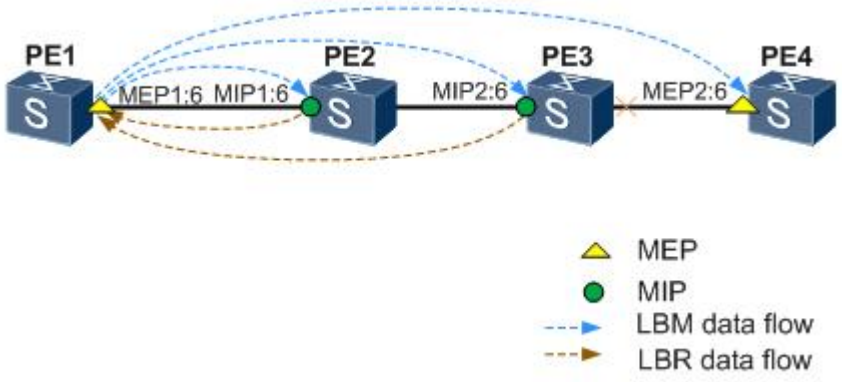

As shown in Figure 3-4, CFM is configured to monitor a path between PE1 (MEP1) and PE4 (MEP2). The MD level of these MEPs is 6. A MIP with the level of 6 is configured on PE2 and PE3. If a fault on a link between PE1 and PE4 is detected by CFM or CCM, 802.1ag MAC ping can be used to locate the fault.

3.2.3 Linktrace

Linktrace is also called 802.1ag MAC trace. Similar to IP traceroute, linktrace identifies a path between two MEPs.

A MEP initiates an 802.1ag MAC trace test to monitor reachability of a MEP or MIP. The initiator and destination node must have the same level. A source MEP constructs and sends an LTM to a destination MEP or MIP. After receiving the LTM, each MIP forwards the LTM and replies with an LTR. Upon receiving the LTM, the destination MEP replies with an LTR message and does not forward the LTM. The MIPs add information about the local device in outgoing LTMs, so the RMEP can obtain information about the entire path. LTMs are multicast packets and LTR messages are unicast packets.

Figure 3-5 Linktrace networking

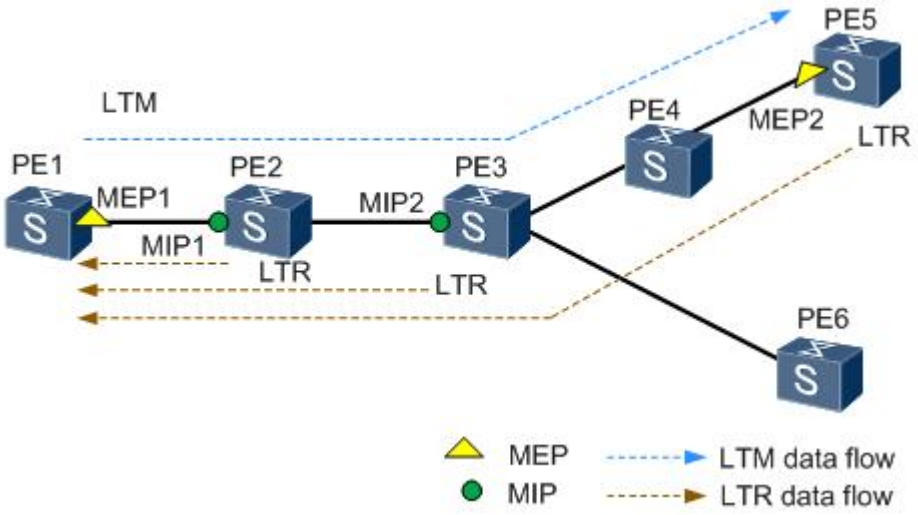

The following example illustrates implementation of linktrace on the network shown in Figure 3-5.

1. MEP1 sends an LTM to MEP2. The LTM carries a time to live (TTL) value and the MAC address of destination MEP2.
2. After receiving the LTM, MIP1 forwards the LTM if the TTL value minus 1 is not 0. MIP1 then replies with an LTR message to MEP1. The LTR carries information about the forwarding path and the TTL field carried in the received LTM.
3. After the LTM reaches MIP2 and MEP2, the process described on MIP1 is repeated for MIP2 and MEP2. In addition, MEP2 finds that its MAC address is the destination address carried in the LTM and therefore does not forward the LTM.
4. The LTRs from MIP1, MIP2, and MEP2 provide MEP1 with information about the forwarding path between MEP1 and MEP2.

If a fault occurs on the path between MEP1 and MEP2, MEP2 or a MIP cannot receive the LTM or reply with an LTR message. MEP1 can locate the faulty node based on such a response failure. For example, if the link between MEP1 and MIP2 works properly, but the link between MIP2 and MEP2 is faulty, MEP1 can receive LTR messages from MIP1 and MIP2 but fails to receive a reply from MEP2. MEP1 then considers the path between MIP2 and MEP2 faulty.

4 Y.1731 Implementation

4.1 Basic Concepts

Y.1731 is Ethernet OAM defined by the ITU-T, and CFM is OAM defined by the IEEE. The ITU-T and IEEE define the OAM protocol together. The two protocols use the same packet format. In addition to fault management of CC, LB, and LT based on CFM, Y.1731 provides performance management.

4.1.1 ME and MEG

A maintenance entity (ME) refers to an entity that requires management and represents the relationship between two maintenance entity group points (MEPs). The ME can be considered as one point-to-point connection. MEs can be nested but cannot be overlapped.

An ME group (MEG) contains one ME or a group of MEs. MEs in an MEG have the same MEG level. The MEG is similar to the MA in 802.1ag.

4.1.2 MEP and MIP

A MEP is an edge node of an MEG, and a MIP is the intermediate node of an MEG. The MEP and MIP of Y.1731 are similar to those of 802.1ag. A MEP sends and terminates OAMPDUs in the same MEG; a MIP only responds to some OAMPDUs, for example, the MIP used as the intermediate node of linktrace replies with LTR messages.

4.1.3 MEG Level

The MEG level is similar to the MD level in 802.1ag. When MEGs are nested, OAMPDUs in a low-level MEG cannot enter a high-level MEG.

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.

4.1.4 Y.1731 Protocol Packet

The Y.1731 packet format is similar to the 802.1ag packet format. Table 4-1 describes the values of the OpCode field.

Table 4-1 Values of the OpCode field

OpCode Value	OAMPDU Type	Description
OpCode values that are defined in the ITU-T Y.1731 and IEEE 802.1		
1	CCM	ITU-T Y.1731 provides the same functions of IEEE 802.1ag.
2	LBR	
3	LBM	
4	LTR	
5	LTM	
0, 6-31, 64-255	Reserved for IEEE 802.1	-
OpCode values that are defined in only the ITU-T Y.1731		
33	Alarm Indication Signal (AIS)	Is used to suppress alarms.
35	Locked Signal (LCK)	Is used to notify administrative locking and subsequent data service interruption.
37	Test Signal (TST)	Is used to perform unidirectional online or offline diagnosis and tests such as bandwidth throughput check, frame loss check, and bit error code check.
39	Automatic Protection Switching (APS)	Is used to control protection switching and enhance reliability.
41	Maintenance Communication Channel (MCC)	Is used for remote management of MEPs.
43	Loss Measurement Message (LMM)	-
42	Loss Measurement Reply (LMR)	-
45	One-way Delay Measurement (1DM)	-
47	Delay Measurement Message (DMM)	-
46	Delay Measurement Reply (DMR)	-
49	Experimental OAM Message (EXM)	-
48	Experimental OAM Reply (EXR)	-

OpCode Value	OAMPDU Type	Description
51	Vendor Specific OAM Message (VSM)	-
50	Vendor Specific OAM Reply (VSR)	-
32, 34, 36, 38, 44, 52-63	Reserved for ITU-T Y.1731	-

4.2 Working Mechanism

Y.1731 covers items defined in IEEE 802.1ag and provides additional OAM messages for fault management and performance monitoring. Fault management includes Alarm Indication Signal (AIS), Remote Defect Indication (RDI), Locked Signal (LCK), Test Signal (TST), Automatic Protection Switching (APS), Maintenance Communication Channel (MCC), Experimental (EXP) OAM, and Vendor Specific (VSP) OAM. Performance monitoring includes frame loss measurement (LM) and delay measurement (DM).

Y.1731 provides the same fault management function as CFM, and the fault management function is not provided here. Performance management includes frame loss measurement, frame delay measurement, frame jitter measurement, and throughput measurement. Y.1731 is a supplement of 802.1ag, and some functions of Y.1731 are not implemented by vendors.Huawei switches provide one-way and two-way frame delay measurement and AIS.

4.2.1 One-Way Frame Delay Measurement

A MEP sends a 1DM frame carrying one-way ETH-DM information to its RMEP. After receiving this message, the RMEP measures the one-way frame delay or delay variation.

One-way frame delay measurement can be implemented only after a MEP synchronizes the time with its RMEP. The delay variation can be measured regardless of whether the MEP synchronizes the time with its RMEP. If a MEP synchronizes its time with its RMEP, the one-way frame delay and delay variation can be measured. If the time is not synchronized, only the one-way delay variation can be measured.

Figure 4-1 One-way frame delay measurement

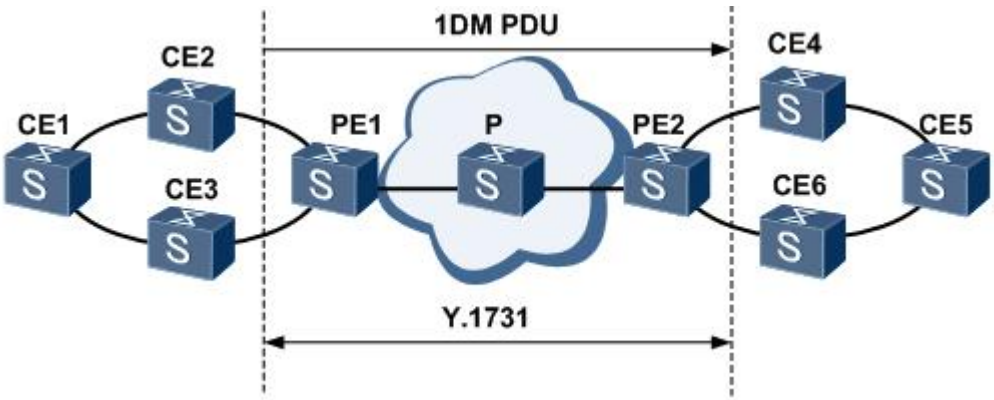

One-way frame delay measurement is implemented on an end-to-end link between a local MEP and its RMEP by exchanging 1DM frames. After one-way frame delay measurement is configured, a MEP periodically sends 1DM frames carrying TxTimeStampf (the time when the 1DM frame was sent). After receiving a 1DM frame, the RMEP parses TxTimeStampf and compares this value with RxTimef (the time when the 1DM frame was received). The RMEP calculates the one-way frame delay based on these values using the following formula:

Frame delay = RxTimef – TxTimeStampf

The frame delay can be used to measure the delay variation.

A delay variation is an absolute difference between two delays that were calculated currently and last time.

4.2.2 Two-Way Frame Delay Measurement

A MEP sends a Delay Measurement Message (DMM) carrying an ETH-DM request to its RMEP. After receiving the DMM, the RMEP sends a Delay Measurement Reply (DMR) carrying an ETH-DM response to the MEP.

Figure 4-2 Two-way frame delay measurement

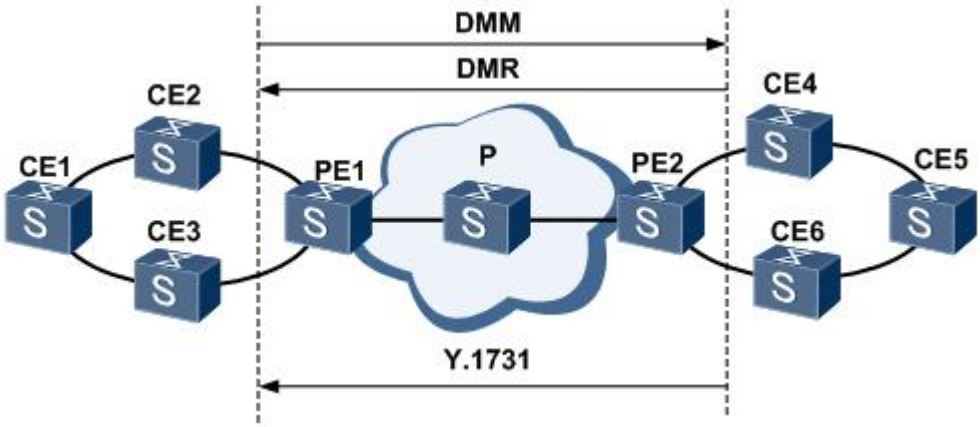

Two-way frame delay measurement is performed by a local MEP to send a DMM to its RMEP and then receive a DMR from the RMEP. After the two-way frame delay measurement is configured, a MEP periodically sends DMMs carrying TxTimeStampf (the time when the DMM was sent). After receiving the DMM, the RMEP replies with a DMR. The DMR carries RxTimeStampf (the time when the DMM was received) and TxTimeStampb (the time when the DMR was sent). The value in every field of the DMM is copied to the DMR except that the source and destination MAC addresses were interchanged. Upon receiving the DMR, the MEP calculates the two-way frame delay by using the following formula:

Frame delay = (RxTimeb - TxTimeStampf) - (TxTimeStampb - RxTimeStampf)

The frame delay can be used to measure the delay variation.

A delay variation is an absolute difference between two delays that were calculated currently and last time.

4.2.3 AIS

AIS is used to transmit fault information.

As shown in Figure 4-3, a MEP is configured in MD1 with the level of 6 on access interfaces of CE1 and CE2 access interfaces on the user network. CE1 and CE2 belong to user networks and have low requirements for fault detection. A MEP is configured in MD2 with the level of 3 on PE1 and PE2. PE1 and PE2 belong to the carrier network and have high requirements for fault detection.

Figure 4-3 Implementation of AIS

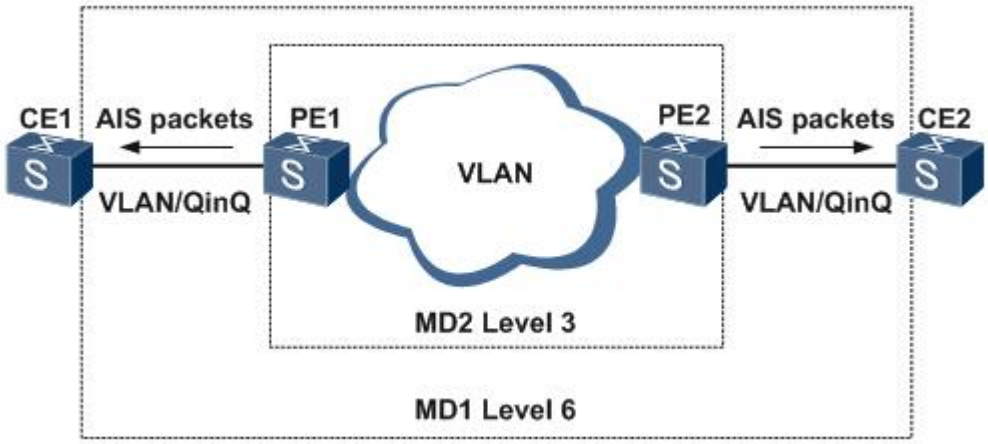

If CFM detects a fault in the link between AIS-enabled PEs, CFM sends AIS packet data units (PDUs) to CEs. After receiving the AIS PDUs, the CEs suppress alarms, minimizing the impact of a lot of alarms on a network management system (NMS).

After the link between the PEs recovers, the PEs stop sending AIS PDUs. CEs do not receive AIS PDUs during a period of 3.5 times the sending interval of AIS PDUs. Therefore, the alarm suppression function on CEs will be canceled.

5 Ethernet OAM Association

5.1 Principle of Ethernet OAM Association

As networks develop quickly, more and more IP networks carry multiple services such as voice and video services. These services pose high requirements on network reliability. Networks need to provide fast fault detection and processing. Link detection protocols are usually deployed on a network to detect link connectivity and faults between devices. A single fault detection protocol cannot detect all faults in all links on a complex network. Various detection techniques are required to rapidly detect and report faults according to actual networking.

Ethernet OAM can detect faults in Ethernet links and advertises fault information to interfaces or other protocol modules. Ethernet OAM association is implemented by the OAM module, application module, and detection module, as shown in Figure 5-1. The detection module can be EFM, CFM, BFD, MPLS OAM, or interface. The application module can be ERPS or SEP. The OAM module associates one module with another. The detection module monitors the link status and network performance. When detecting a fault, the detection module instructs the OAM module to notify the application module or another detection module of the fault. After receiving the notification, the application or detection module takes measures to prevent a communication interruption or service quality deterioration.

Figure 5-1 Ethernet OAM association

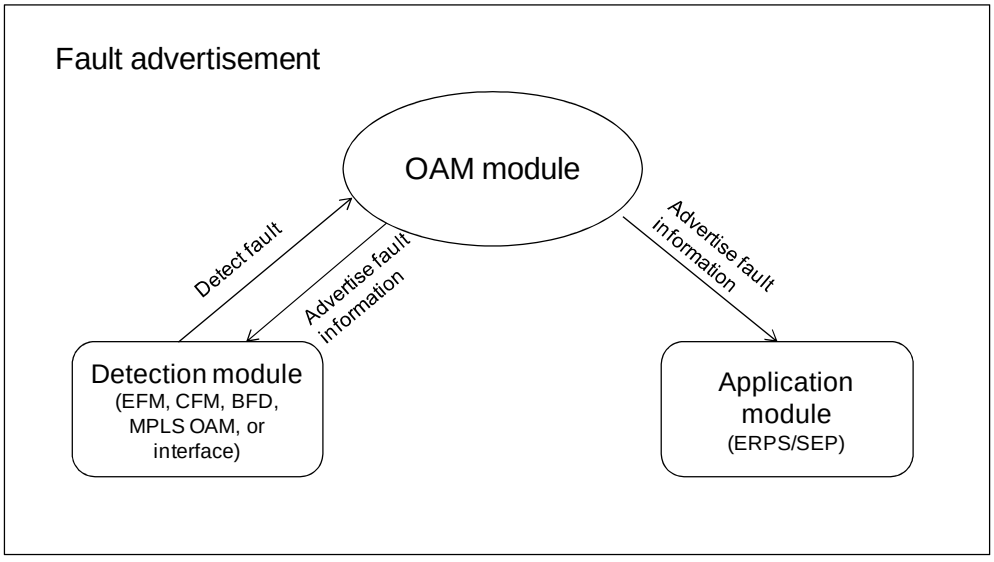

Protocols of the same type or different types of the detection module can be associated. Fault information can be transmitted unidirectionally or bidirectionally. A protocol of the detection module is associated with the application module to accelerate fault detection, implement fast convergence, and reduce service interruption. Protocols or interfaces of the detection module are associated to transmit fault information on different links, so devices along multiple links can respond to and handle faults and switch services to the standby path. As shown in Figure 5-2, the normal service forwarding path is UPE1 -> PE2 -> PE4 -> PE6 -> PE8. EFM is deployed between CE1 and UPE1, CFM1 is deployed between UPE1 and PE4, and CFM2 is deployed between PE4 and PE8. If a Layer 3 network is deployed between PE4 and PE8, BFD is used to detect faults on the link. Association between CFM1 and EFM, between CFM1 and Port1, between CFM1 and CFM2/BFD, and between CFM2/BFD and Port2 is deployed. When CFM1 detects a fault, CFM1 advertises fault information to EFM and CFM2/BFD. Then UPE1 and PE8 switch services to the standby path UPE1 -> PE1 -> PE3 -> PE5 -> PE7, preventing service interruption.

Figure 5-2 Associating multiple protocols to transmit fault information

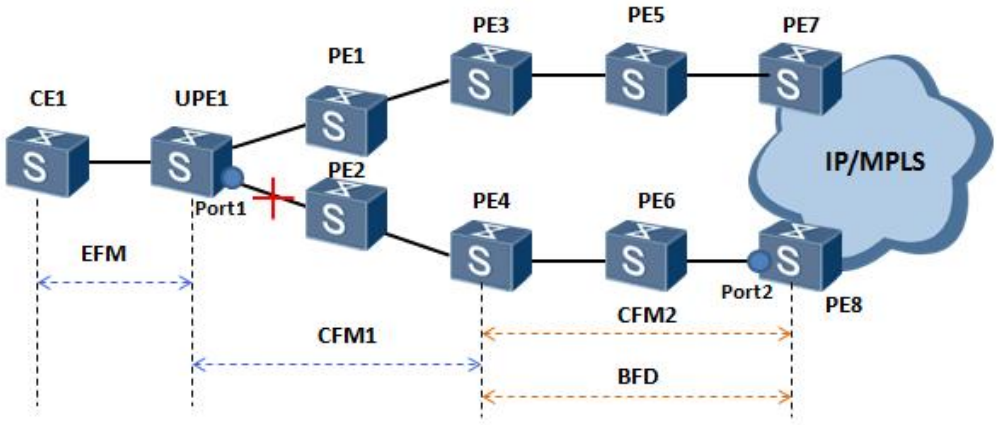

5.1.1 Association Between EFM Modules

When the EFM module at one side detects a fault, the OAM module sends the fault message to the associated EFM module at the other side. The fault report mode can be the following:

- The EFM module at one side sends fault messages to the EFM module at the other side.
- The EFM modules at both sides send fault messages to each other.

5.1.2 Association Between EFM and an Interface

When an EFM-enabled interface detects a link connectivity fault through EFM, the associated interface is configured in Down state. EFM can be associated with an interface of the detected link or non-detected link.

Association between EFM and an interface of the detected link: When the EFM module detects a link fault, the interface is configured in ETHOAM Down state. Only EFM OAMPDUs are transmitted and other service packets or flows are blocked, speeding up route convergence. In addition, services can be switched from the active path to the standby path.

Association between EFM and an interface of the non-detected link: When the EFM module detects a fault, the physical status of the interface bound to the EFM module becomes TRIGGER DOWN (3AH). The OAM module sends a fault message to the bound EFM module and then sends EFM OAMPDUs to the remote device to notify the fault. The fault report mode can be the following:

- When association between EFM and an interface of the detected link is used, only the EFM module can send fault messages to the bound interface.
- When association between EFM and an interface of the non-detected link is used, the EFM module can send fault messages to the bound interface.
- When association between EFM and an interface of the non-detected link is used, the interface can send fault messages to the bound EFM module.
- When association between EFM and an interface of the non-detected link is used, the EFM module and an interface can send fault messages to each other.

5.1.3 Association Between EFM and CFM

When CFM detects a fault in an MA, the OAM module sends a fault message to the associated EFM module. When EFM detects a fault, the OAM module sends a fault message to the associated CFM module in the MA. The fault report mode can be the following:

- The CFM module sends fault messages to the EFM module.
- The EFM module sends fault messages to the CFM module.
- The CFM module and the EFM module send fault messages to each other.

5.1.4 Association Between EFM and BFD

When EFM detects a fault, the OAM module sends a fault message to the associated BFD session. When a BFD session detects a fault, the BFD session sends a fault message to the associated EFM module. The fault report mode can be the following:

- The EFM module sends fault messages to a BFD session.
- A BFD session sends fault messages to the EFM module.
- The EFM module and BFD session send fault messages to each other.

5.1.5 Association Between EFM and MPLS OAM

When EFM detects a fault, the OAM module sends a fault message to MPLS OAM. When MPLS OAM detects a fault, the OAM module sends a fault message to the associated EFM module. The fault report mode can be the following:

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

5.1.6 Association Between CFM and CFM

When the CFM module detects a fault in an MA, the OAM module reports the fault to the bound MA at the other side. The fault report mode can be the following:

- The CFM module at one side sends fault messages to the CFM module at the other side.
- The CFM modules at both sides send fault messages to each other.

5.1.7 Association Between CFM and an Interface

When CFM detects a connectivity fault between a MEP in an MA and its RMEP, the associated interface is configured in Down state. CFM can be associated with an interface of the detected link or non-detected link.

Association between CFM and an interface of the detected link: When CFM detects a link fault, the interface where the MEP is configured is shut down and enable after 7s so that other modules can detect the fault.

Association between CFM and an interface of the non-detected link: When CFM detects a fault, the physical status of the interface bound to the CFM module goes Down, speeding up route convergence of the upper-layer protocol. When the interface associated with the CFM module goes Down, the OAM module sends a fault message to the associated CFM module, and sends CFM protocol packets to the remote device to notify the fault. The fault report mode can be the following:

- When association between CFM and an interface of the detected link is used, only the CFM module can send fault messages to the bound interface.
- When association between CFM and an interface of the non-detected link is used, the CFM module can send fault messages to the bound interface.
- When association between CFM and an interface of the non-detected link is used, the interface can send fault messages to the bound CFM module.
- When association between CFM and an interface of the non-detected link is used, the CFM module and an interface can send fault messages to each other.

5.1.8 Association Between CFM and BFD

When CFM detects a fault in an MA, the OAM module sends a fault message to the associated BFD session. When a BFD session detects a fault, the BFD session sends a fault message to the bound MA. The fault report mode can be the following:

- The CFM module sends fault messages to a BFD session.
- A BFD session sends fault messages to the CFM module.
- The CFM module and BFD session send fault messages to each other.

5.1.9 Association Between CFM and MPLS OAM

When CFM detects a fault in an MA, the OAM module sends a fault message to MPLS OAM. When MPLS OAM detects a fault, the OAM module sends a fault message to the CFM module in the MA. The fault report mode can be the following:

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

5.1.10 Association Between CFM and SEP

SEP works at the access or aggregation layer. To ensure that the network running SEP immediately detects the network topology change, deploy association between CFM and SEP.

When CFM detects a fault on a link, the edge device notifies the OAM module of the fault by sending a CFM protocol packet. Then the SEP status of the interface associated with CFM on the edge device changes to Down. When the interface bound to CFM on the edge device becomes Down, an interface of the remote device in the SEP segment needs to send Flush-FDB packets to report the topology change to other nodes in the SEP segment. After devices in the SEP segment receive Flush-FDB packets, the blocked port in the SEP segment switches to the Forwarding state and sends Flush-FDB packets to trigger other nodes in the SEP segment to update their MAC address table and ARP table. In this manner, faults can be rapidly detected and the reliability of service transmission is thus ensured.

5.1.11 Association Between CFM and ERPS

After an interface in an ERPS ring is associated with CFM, fault detection is accelerated. After CFM detects a fault, CFM notifies ERPS of updating the MAC address table and ARP table, implementing fast convergence and reducing traffic interruption. ERPS cannot automatically detect link faults. When there are relay transmission devices in an ERPS ring, ERPS cannot detect whether faults on relay devices cause slow convergence and traffic interruption. Association between CFM and ERPS can solve this problem.

6 Implementation on Huawei Products

6.1 Hardware-based 3.3-ms Fault Detection

Huawei switches provide industry-leading hardware-based link detection technology. The switch enabled with Ethernet OAM sends detection packets at an interval of 3.3 ms, which is the fastest detection time in industry. When a fault occurs on a Layer 2 link, services can be switched within 50 ms. Users do not detect network faults. Independent hardware is used to send detection packets, so the CPU does not detect the event and no CPU resource is consumed. When detection packets are sent to detect faults, this process is not affected by the CPU. The reliability is ensured in case of heavy traffic. In addition, Huawei switches also provide hardware-based 3.3-ms BFD to ensure that services are switched within 50 ms when a fault occurs on a Layer 3 link. OAM and BFD can be associated and transmit fault messages to each other, which implements fast E2E fault detection. In addition, Layer 2 and Layer 3 services can be rapidly switched.

6.2 Various OAM Association Technologies

Link detection protocols are usually deployed on a network to detect link connectivity between neighboring devices. A single fault detection protocol can hardly detect all faults in all links on a complex network, so various detection techniques are required to rapidly detect and report faults according to actual networking. EFM and CFM can be associated with an interface, BFD, MPLS OAM, and ERPS/SEP/Smart Link/RRPP. In addition, EFM or CFM can be associated with itself. The association function implements fast advertisement of fault information and ensures E2E service reliability.

7 Application Scenarios

7.1 Associating EFM with an Interface to Enhance Network Reliability

As shown in Figure 7-1, switches and routers constitute an RRPP ring. Routers are configured with different VPLS VSIs to transparently transmit RRPP packets and service data packets. EFM is configured between switches and between switches and routers; EFM is associated with an interface. When the link quality is low, EFM may detect a fault after the timeout. Then EFM instructs the associated interface to enter the Down state. In this situation, RRPP flapping may occur frequently, causing temporary loops. You can set the value of EFM faulty-state holdup timer so that the EFM status does not frequently alternate between Up and Down in a short period of time in the case of low link quality. The service stability is therefore improved.

Figure 7-1 Networking of association between EFM and an interface

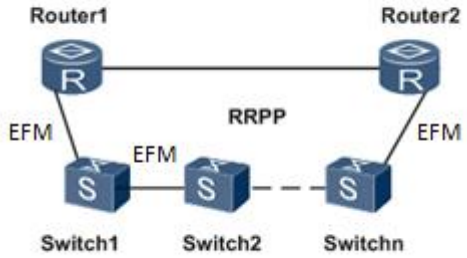

You can also configure association of Errored Symbol Event or Errored Frame event. When the number of errored frames, errored codes, or errored frame seconds detected by the interface reaches or exceeds the threshold within a period of time, the link is considered unavailable or its quality is low. Then the administrative status of the interface becomes Down and services are switched to the standby link. Set the recovery time of the interface as needed. After the interface is restored, use EFM to detect faults.

7.2 Associating CFM with SEP in Dual-Homing Networking

As shown in Figure 7-2, LSW1 to LSW5 run SEP to implement redundancy at the access layer. Association between SEP and CFM is deployed on the edge device LSW1 of the SEP segment. When CFM detects a fault at the aggregation layer, LSW1 encapsulates a fault message into a CFM protocol packet and sends it to the OAM module to trigger the physical status of the interface bound to CFM to Down. Then services are switched between links.

Figure 7-2 Networking of association between SEP and CFM

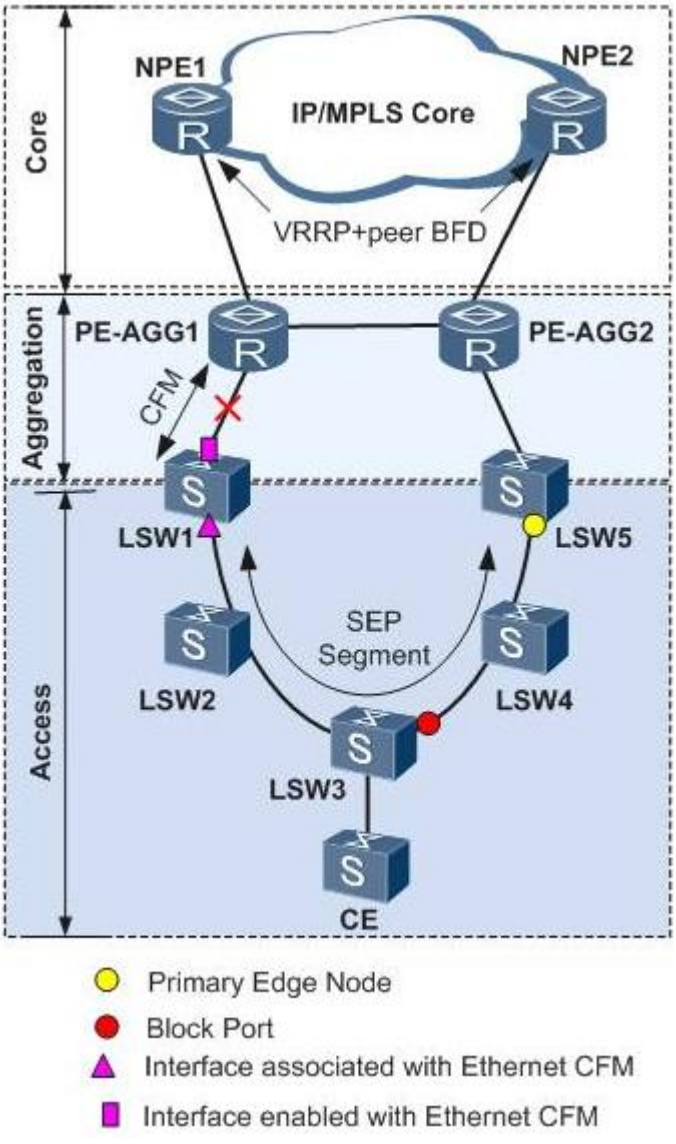

7.3 E2E Multi-Link Detection and Protection

As shown in Figure 7-3, EFM is deployed at the access layer. If L2VPN (VPLS/VLL) is deployed at the core layer and CFM is deployed between UPE1 and UPE2, VPLS or VLL can transparently transmit CFM packets. If the core layer uses native IP or L3VPN, BFD or MPLS OAM can be deployed at the core layer, CFM can be deployed at the aggregation layer, and CFM is associated with BFD/MPLS OAM and EFM. An E2E multi-link detection and protection solution is therefore deployed.

Figure 7-3 Networking of E2E multi-link detection

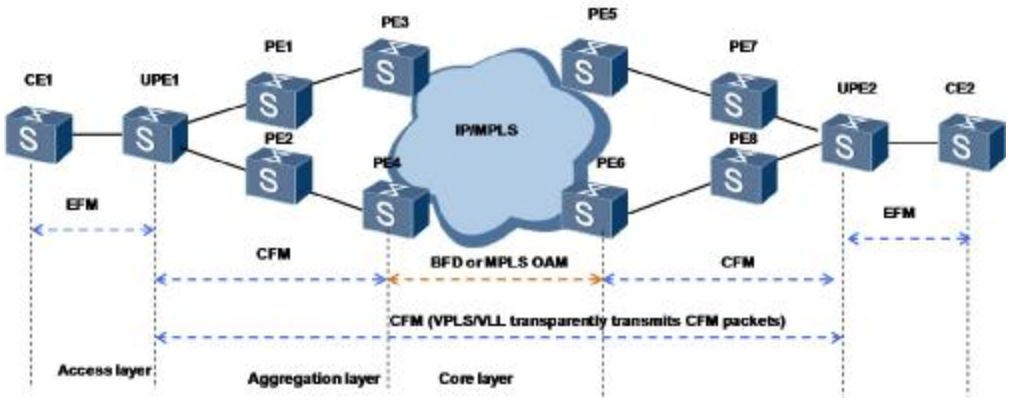

8 Appendix

8.1 References

Standard Number	Name	Description
IEEE Std 802.3ah-2004	Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks	Functions: 1. Discovers OAM capabilities of the remote device. 2. Monitors the link status and generates an event notification when the error threshold is reached. 3. Instructs the remote end to receive faults through RDI. 4. Supports remote loopback. Loopback needs to be performed on the remote end and statistics can be collected through testing.
IEEE Std 802.1ag-2007	IEEE Standard for Local and metropolitan area networks-Virtual Bridged Local Area Networks Amendment 5:Connectivity Fault Management	Functions: 1. Defines MDs of different levels. An MD can contain multiple MIPs and MEPs. 2. Defines packets of CC, LB, and LT based on the standard frame format to implement OAM functions. 3. Provides OAM capabilities based on E2E Ethernet virtual connections.
ITU-T Y.1731	OAM functions and mechanisms for Ethernet based networks	Provides the following functions based on CFM: 1. Offers ETH-LCK. 2. Defines multicast loopback packets. 3. Tests and maintains communication channels. 4. Measures the performance including the delay and packet loss.

Standard Number	Name	Description
MEF E-LMI	Ethernet Local Management Interface	E-LMI is an asymmetrical protocol and defines OAM management from a UPE to a CE. It can be only used on UNIs. E-LMI defines the process and message format. Through E-LMI, the UPE can exchange multiple types of information with a CE, including the EVC status, remote UNI status, mapping between VLANs and EVCs on the CE, and bandwidth profile.

8.2 Acronym and Abbreviation

Acronym	Full Name
IEEE	Institute of Electrical and Electronics Engineers
ITU-T	International Telecommunication Union-Telecommunication Standardization Sector
MEF	Metro Ethernet Forum
UNI	User Networks Interface
EVC	Ethernet Virtual Connection
SLA	Service Level Agreement
CCM	Continuity Check Message
CFM	Connectivity Fault Management
EFM	Ethernet of First Mile
LB	Loopback
LT	Linktrace
MA	Maintenance Association
MD	Maintenance Domain
MEP	Maintenance End Point
MIP	Maintenance Intermediate Point
RMEP	Remote Maintenance association End Point
OAM	Operation, Administration and Maintenance
OAMPDU	OAM Packet Data Unit
RDI	Remote Default Indication

Acronym	Full Name
ME	Maintenance Entity
MEG	Maintenance Entity Group