OAM and Its Performance Monitoring Mechanisms for Carrier Ethernet Transport Networks

Jeong-dong Ryoo, Jongtae Song, Jaewoo Park, and Bheom-Soon Joo, ETRI

ABSTRACT

Ethernet technology is rapidly gaining importance as it becomes a dominant solution for a converged transport network. Ethernet OAM features defined in standards provide a means of performance improvement to meet carrier-class transport network requirements. This article outlines Ethernet OAM functions and mechanisms, and explains how its performance monitoring schemes work. In addition, this article introduces open issues and their potential solutions in the performance monitoring of Ethernet OAM for the next phase of standardization.

INTRODUCTION

Based on the forecast of exponential growth of data traffic, which is mainly IP packet traffic, network operators have built overlay networks for data transport. The use of overlay networks is preferred in this early stage to cope with the demand on data transport, as a new service can be introduced without affecting legacy infrastructure. Since there are multiple networks to run in parallel, this solution results in bigger capital expenditures (CAPEX) and operational expenditures (OPEX). To address the issues in the overla approach, it is generally considered that the next-generation transport network will be a pure packet-based network. Ethernet is one of the strong candidates for the future, because of its fast data transfer and simple deployment and, more important, great compatibility with IP traffic. This powerful and low-cost networking technology is now gaining strong attention for applications in the next-generation transport network.

It is widely recognized that operations, administration, and maintenance (OAM) are important in carrier transport networks to ease network operation, verify network performance, and deliver availability and reliability objectives to support the requirements of a service level agreement (SLA). As Ethernet technology seeks its way to public carrier transport networks, Ethernet has driven the need for a new and powerful set of OAM tools. There are also many other motivations for the new Ethernet OAM, such as:

- In order to manage and troubleshoot layer 2 Ethernet service, overlaying the IP infrastructure is a burden.
- Any lower-layer OAM mechanism cannot act as a substitute for Ethernet OAM. A new Ethernet OAM can deal only with situations relevant to Ethernet, the nature of which is both connectionless and multipoint connected.
- Ethernet in the first mile (EFM) OAM is good for single-link connections, but cannot monitor across Ethernet virtual connections (EVCs).
- Independent of underlying technologies, such as native Ethernet, Ethernet over synchronous optical network (SONET), Ethernet over asynchronous transer mode (ATM), Ethernet over MPLS, and Ethernet over resilient packet ring (RPR), end-to-end Ethernet services have to be monitored across diverse networks.
- There is strong demand for multipoint-to-multipoint Ethernet service, which must be covered in a new Ethernet OAM.
- There is a need to determine availability and network performance to ensure that customers are charged properly for subscribed Ethernet services.

Standards bodies like the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T), IEEE, and Metro Ethernet Forum (MEF) have been working closely to drive consistent recommendations and standards for Ethernet OAM, and the first phase of the work is completed in defining Ethernet OAM functions and mechanisms with efforts in both IEEE 802.1ag under the name "Connectivity Fault Management" [1] and ITU-T SG13 under the name "Y.1731 — OAM Functions and Mechanisms for Ethernet Based Networks" [2]. As most of these standards

 0163-6804/08/$25.00

Ethernet OAM operates with the notion of hierarchical maintenance domains. In other words, multiple levels of maintenance domains can be managed with a single OAM mechanism. The maintenance domain is called Maintenance Entity Group.

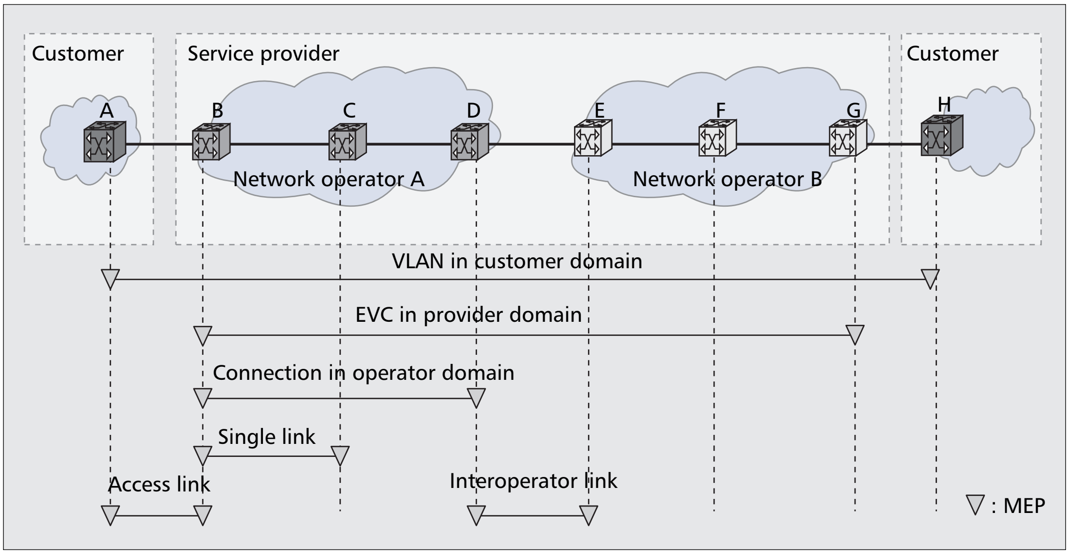

Figure 1. *An Ethernet OAM reference model.*

mature, a new generation of Ethernet networking is imminently anticipated.

In this article we provide an overview of relevant standards and summarize key protocol aspects. We mainly focus on the OAM functions and mechanisms from the ITU-T perspective, as the functionalities defined by the ITU-T forms a superset of those in other standards. This article also introduces open issues and their proposed solutions in the performance monitoring of Ethernet for the next stage of standardization.

ETHERNET OAM

Ethernet OAM operates with the notion of hierarchical maintenance domains. In other words, multiple levels of maintenance domains can be managed with a single OAM mechanism. The maintenance domain is called a maintenance entity group (MEG). According to Y.1731 and IEEE 802.1ag, a maximum of eight levels are possible. Figure 1 serves as a reference model to be used for this article, and shows typical maintenance domains in different levels. The boxes lettered from A to H represent Ethernet switching equipment. There can be multiple organizations involved in Ethernet services: customers, network operators, and service providers. Each organization, including the customer, can run Ethernet OAM independently at its own level to manage and monitor the OAM domains for which it is responsible. The OAM frames belonging to higher levels are transparently forwarded by lower-level switches. In this example, node D in network operator A transparently passes Ethernet OAM frames from the customer (node A–node H) and service provider (node B–node G), and the customer never sees the network operator's Ethernet OAM frames.

Ethernet OAM has two main features: fault management and performance monitoring. Fault management allows detection, verification, localization, and notification of different defect conditions. The performance monitoring allows measurement of different performance parameters, such as loss, delay, and jitter. Table 1 summarizes the OAM functionalities provided by ITU-T Y.1731 and IEEE 802.1ag. Mechanisms supported by IEEE 802.1ag include continuity check (CC), loopback (LB), link trace (LT), and remote defect indication (RDI, embedded in continuity check message [CCM]). ITU-T Y.1731 supports a rich set of OAM functions as listed in Table 1.

FAULT MANAGEMENT

CCM is arguably the most important of the OAM messages defined in both Y.1731 and 802.1ag. CCM is defined as a proactive OAM message, which means that once started it is automatically generated at a configured rate, while other messages are on demand. The transmission period ranges from 3.33 ms to 10 min. The CCM enables detection of loss of continuity between the endpoints in an OAM domain, called MEG endpoints (MEPs). An MEP declares a loss of continuity when it does not receive the expected CCM for 3.5 times the configured transmission period. The loss of continuity indicates a fabric or link failure between two endpoints in point-to-point connection. CCM also performs various other defect and performance monitoring activities, such as detection of unintended connections between MEPs, RDI, and frame loss measurement.

Similar to *ping* in IP networks, LB is used to verify bidirectional connectivity to a particular MEG intermediate point (MIP) or MEP. After an LB request message (LBM) is sent on demand, an LB reply (LBR) is expected to be received within 5 s. If not, the connectivity to the peer cannot be verified. LBM/LBR can optionally carry test patterns for various diagnostic tests, such as verifying bandwidth throughput and detecting bit errors. The purpose of the test signal (TST) is the same as that of LB, but the TST function is used to perform one-way diagnostic tests.

LT function can be used to identify the path between two OAM entities. This is similar to *traceroute* in IP. Especially within a multipoint-to-multipoint connectivity environment like E-

LAN, LT is quite a useful tool to retrieve the adjacency relationship between an MEP and a remote MEP or MIP.

Alarm indication signal (AIS) messages are used to suppress alarms at client layers from defects discovered at server layers. When an MEP detects a connectivity failure at a level, it starts transmitting periodic AIS frames in the direction away from the detected failure at the next higher level. For example, in Fig. 1, when the MEP for the "connection in operator domain" level in node D detects the loss of continuity in the direction from node B to node D, node D transmits AIS frames at the "EVC in provider domain" level in the direction to node G. As node G receives AIS from node D, it also generates AIS for node H. Due to independent restoration capabilities provided within the Spanning Tree Protocol (STP), AIS is not applied in STP environments, as STP has its own mechanism for this purpose.

The locked signal (LCK) message is used to inform MEPs of intentional diagnostic actions, enabling client MEPs to differentiate between a defect condition and an administrative locking action at the server layer MEP.

Automatic protection switching (APS) is used to control linear protection switching operations. Two disjointed transport entities (e.g., two different VLAN IDs) are used as working and protection transport entities. Being a 1-phase APS protocol, Ethernet linear protection switching requires only a single information exchange between two MEPs to complete a protection switching. Therefore, faster switching time can be achieved than two- or three-phase APS protocols. Three protection switching architectures are supported: 1+1 unidirectional, 1+1 bidirectional and 1:1 bidirectional. In unidirectional protection switching, each direction is switched independently. As for bidirectional protection switching, a failure in one direction causes protection switching on both directions. In the case of 1+1 uni-/bidirectional protection switching, one MEP duplicates data frames and sends through both working and protection transport entities. 1:1 bidirectional protection switching allows only one transport entity to deliver data frames. At the other MEP, data frames from only one transport entity are selected in any cases. Revertive and non-revertive operations are also defined. Ethernet linear protection switching is detailed in ITU-T G.8031 [3]

The maintenance communication channel (MCC) provides a communication channel between a pair of MEPs in order to make remote maintenance possible. Experimental OAM is for trying out new OAM functionality within an administrative domain on a temporary basis. Vendor-specific OAM is used for vendor-specific OAM functionality and may be used by a vendor across its equipment.

Performance Monitoring

The current Recommendation Y.1731 specifies four performance parameters: frame loss ratio, frame delay, frame delay variation, and throughput. Frame loss ratio is the ratio of the number of service frames not delivered divided by the total number of service frames during a given time interval. The method to measure the number of lost frames is described later in this section. Frame delay is measured with either a 1DM message for one-way delay measurement or DMM-DMR messages for two-way delay measurement, and frame delay variation is calculated from the values obtained by frame delay measurement. Throughput is defined as the maximum rate at which no frame is dropped, and its measurement can be performed with LBM-LBR or TST messages.

Frame loss measurement can be performed in two ways: dual-ended and single-ended. Dual-ended frame loss measurement utilizes CCM frames, and is used as proactive OAM. For the single-ended frame loss measurement, LMM-LMR frames are used, and the measurement is performed on demand. The mechanisms for both measurement schemes are similar; we explain the dual-ended frame loss measurement scheme. For the dual-ended frame loss measurement in the case of point-to-point ME, each of two peering MEPs maintains two local counters: TxFCl (for in-profile data frames transmitted toward a peer MEP) and RxFCl (counter for in-profile data frames received from a peer MEP). Each of two peering MEPs periodically transmits CCM frames with the following information related to frame loss measurement:

Function	OAM message	ITU-T Y.1731	IEEE 802.1ag
Continuity check	CCM	Yes	Yes
Loopback	LBM-LBR	Yes	Yes
Link trace	LTM-LTR	Yes	Yes
Alarm indication signal	AIS	Yes	No
Remote defect indication	CCM	Yes	Yes
Lock signal	LCK	Yes	No
Test signal	TST	Yes	No
Automatic protection switching	APS	Yes	No
Maintenance communication channel	MCC	Yes	No
Experimental OAM	EXM-EXR	Yes	No
Vendor-specific OAM	VSM-VSR	Yes	No
Frame loss measurement: dual-ended	CCM	Yes	No
Frame loss measurement: single-ended	LMM-LMR	Yes	No
Frame delay measurement: one-way	1DM	Yes	No
Frame delay measurement: two-way	DMM-DMR	Yes	No
Throughput measurement	LBM-LBR or TST	Yes	No

Table 1. *Ethernet OAM functions provided by ITU-T Y.1731 and IEEE 802.1ag.*

As end-to-end Ethernet services can be provided over various underlying technologies, including SDH, RPR, MPLS, T-MPLS, ATM, etc, Ethernet OAM functions can operate independently with the OAM functions of those lower layer networks.

- TxFCf: Value of TxFCl at the time of CCM frame transmission
- RxFCb: Value of RxFCl at the time of the last CCM frame reception
- TxFCb: Value of TxFCf in the last received CCM frame

Each MEP independently calculates the number of lost frames in both directions between two consecutive arrivals of CCM frames, say at t_p (the reception time of the previous frame) and t_c (the reception time of the current frame). The frame losses in the direction away from the calculating MEP (far-end) and in the direction toward the calculating MEP (near-end) are expressed as

$$\text{Frame Loss}_{\text{far-end}} = |\text{TxFCb}[t_c] - \text{TxFCb}[t_p]| - |\text{RxFCb}[t_c] - \text{RxFCb}[t_p]|$$

and

$$\text{Frame Loss}_{\text{near-end}} = |\text{TxFCf}[t_c] - \text{TxFCf}[t_p]| - |\text{RxFCl}[t_c] - \text{RxFCl}[t_p]|,$$ respectively.

There are two ways defined to measure the frame delay performance: one-way and two-way. In one-way frame delay measurement, an MEP sends the one-way frame delay measurement (1DM) frame with timestamp, and its peering MEP calculates the delay of the 1DM frame based on the timestamp value of the frame and the time at reception of the 1DM frame. In this case the clocks of two MEPs should be synchronized. Without clock synchronization, only frame delay variation measurement can be obtained. Two-way frame delay measurement uses the frame delay measurement request (DMM) and frame delay measurement reply (DMR) frames. An MEP transmits the DMM frame with timestamp, whose value reflects the time at transmission of the DMM frame. When the DMM frame is arrived at peer MEP, a DMR frame is generated with the timestamp value in the DMM frame, and transmitted back to the origin of the DMM frame. The origin of the DMM frame then calculates round-trip frame delay. Optionally, two additional timestamps can be used to take into account the processing time at the receiver of the DMM: timestamp at the time of receiving the DMM frame and timestamp at the time of transmitting the DMR frame. In any case, two-way ETH-DM measurement just produces round-trip delay.

RELATIONSHIP WITH OTHER OAMS

Generally speaking, two approaches are considered when two different OAM mechanisms have to operate together for end-to-end OAM service: interworking and layering. In the interworking approach, the OAM service is transported via the interworking function between two peering networks. The layering approach allows lower-layer networks transport end-to-end OAM frames transparently. The OAM mechanisms for transporting end-to-end layer 2 (e.g., ATM) traffic using multiprotocol label switching (MPLS) networks is presented in [4].

As for the interworking approach, the OAM service interworking between Ethernet and other networks needs to be defined. This involves the conversion of information between Ethernet OAM frames and other OAM packets, and specific processing for each OAM function interworking. The OAM interworking between Ethernet and MPLS, and with transport MPLS (T-MPLS) is described in [5, 6]. In [7] the authors describe the interworking between Ethernet OAM and EFM OAM as well as the principles of Ethernet OAM from the IEEE 802.1ag perspective.

As end-to-end Ethernet services can be provided over various underlying technologies, including SDH, RPR, MPLS, T-MPLS, and ATM, Ethernet OAM functions can operate independently with the OAM functions of those lower-layer networks. For example, Ethernet OAM frames are carried end-to-end, and transported transparently by an MPLS network. When a defect arises in the lower-layer network, the Ethernet layer needs to be notified to recognize the defect and suppress unnecessary alarms. Also, if protection switching functions are enabled for both layers, coordination is needed to let the lower layer act on the failure first. When the lower layer fails to recover within the holdoff time, the Ethernet layer can execute its protection switching function.

UPCOMING ETHERNET OAM ENHANCEMENTS

To enhance and modify the first version of Ethernet OAM (i.e., the current approved Y.1731), several issues are identified for future work. It has been agreed to include the following items in the next version:

- The behavior of LCK is modified to transmit an LCK signal in the directions toward as well as away from its peer MEP.
- CCM and APS frames are counted by two local frame counters in an MEP under some conditions.
- ITU-T SG15 Q9 is currently developing a recommendation for Ethernet ring protection, G.8032, which defines the APS protocol and protection switching mechanisms for Ethernet ring or interconnected rings topology. To support the ring-specific APS channel, a new Ethernet OAM message, called R-APS, will be defined.
- Some clarifications are made to the text.

Other than those items, there are two major study items. The study items are frame loss measurement of multicast related service and priority-based one-way delay measurement. The content of these items were presented at the most recent ITU-T meeting. Attendees agreed on studying these items and collecting feedback from network providers and other standards organizations. In this section we provide detailed information on the two study items, especially their motivations and basic mechanisms.

FRAME LOSS MEASUREMENT FOR MULTICAST SERVICE

Initially, studying multipoint-to-multipoint measurement was agreed on. After investigating the mechanisms for multipoint-to-multipoint frame loss measurement, it was felt that a feasible solution for multipoint-to-multipoint may not be available. There is no obvious way for an ingress MEP to determine from which egress MEP a data frame will leave, or for an egress MEP to determine into which ingress MEP a data frame entered. Frame loss measurement for multicast service traffic is proposed as a complementary solution of multipoint-to-multipoint measurement.

With the popularity of video distribution and other one-to-many applications, multicast Ethernet service is expected to grow dramatically. Since frame loss affects the quality of the video significantly, providers or operators need to measure frame loss performance for a specific multicast service. OAM directly impacts a provider's ability to maintain high SLA for a particular premium multicast service.

Consider a multipoint network in Fig. 2. The multicast service traffic that needs to be monitored flows into the provider network via network element A. The multicast service traffic is delivered to network elements B, D, and F in this example. However, the egress network elements for multicast traffic change over time, and an egress network element depends on the dynamics of a subscriber's join/leave actions. The ingress MEP at (A, G) maintains one local counter for each multicast service being monitored: TxFCl (for a specific multicast service data frame transmitted toward the egress MEPs). In the meantime, each of all other MEPs — (K, F), (L, E), (L, D), (H, B), and (H, C) in the figure — maintains one local counter for each multicast service being monitored: RxFCl (for a specific multicast service data frame received from the ingress MEP of the multicast service). An ingress service frame is duplicated at network elements I and J, which are placed in the core of the provider's network. With the multicast service traffic flow depicted in Fig. 2, only the values of the counters at (B, H), (D, L), and (K, F) points would be increased.

The ingress MEP periodically transmits multipoint frame loss measurement request (say, mLMM) frames to all the other MEPs in its MEG with the following information elements:

- Multicast service ID: ID of the multicast service that needs to be monitored. An example of this value is the IP address (or medium access control [MAC] address) of the multicast group being monitored.
- TxFCf: Value of the local counter TxFCl for a specific multicast service data frame at the time of mLMM frame transmission.

When an mLMM frame is received by an MEP and the receiving MEP is engaged in subscribing the multicast service specified in the multicast service ID of the received mLMM frame, a multipoint frame loss measurement reply (say, mLMR) frame is generated and transmitted to the requesting MEP. An mLMR frame contains the following values:

- Multicast service ID: ID of the multicast service being monitored
- TxFCf: Value of TxFCf copied from the mLMM frame
- RxFCf: Value of the local counter RxFCl at the time of mLMM frame reception

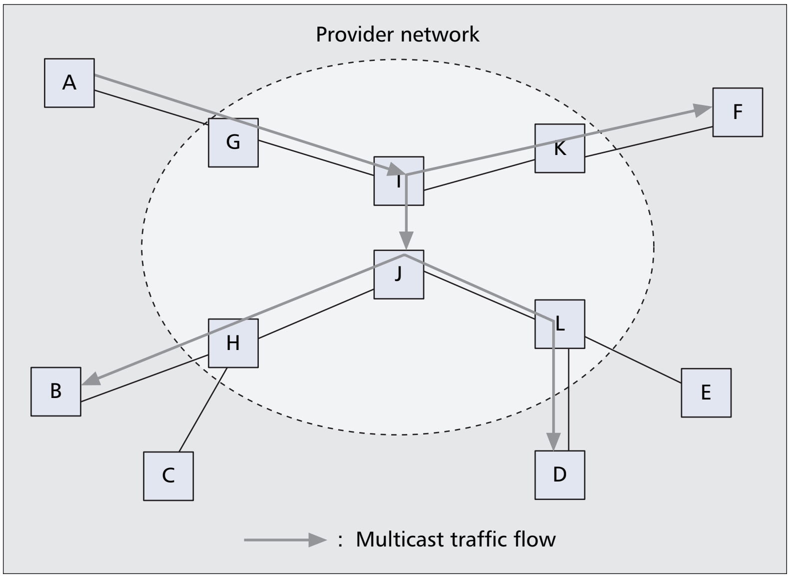

■ **Figure 2.** *Multicast service in a multipoint network.*

When an mLMM frame is received by an MEP but the receiving MEP is not engaged in subscribing to the multicast service specified in multicast service ID of the received mLMM frame, the receiving MEP silently discards the mLMM frame and does not generate any reply frame.

After two consecutive transactions, the MEP receiving mLMR frames (received at t_p and t_c) uses the values of TxFCf and RxFCf fields to make far-end loss measurement for the egress MEP.

When multiple egress MEPs are active on the multicast service, the same number of far-end frame loss values will be obtained. Even though the join/leave actions of subscribers are spontaneous, the scheme presented above does not have to know which network elements are participating in subscription at the time the monitoring is performed. It should also be noted that in the presented scheme an ingress network element is assumed to be known, but with two local counters at each network element this assumption can be relaxed.

Priority-Based One-Way Ethernet Frame Delay Measurement

The one-way Ethernet frame delay measurement scheme in the current Y.1731 [2] cannot be utilized to obtain frame delay when the clocks between two MEPs are not synchronized. As for two-way Ethernet frame delay measurement, the mechanism in the current Y.1731 provides round-trip delay only. However, most Internet applications are client-server whose traffic distribution is usually asymmetric; thus, the round-trip delay does not show us the exact delay performance in a real network.

The service classes of an Ethernet service network are categorized based on the value of the priority field in Ethernet frames. With different priority values of service frames, the switches in the network apply different buffering and scheduling policies for the frames. The highest

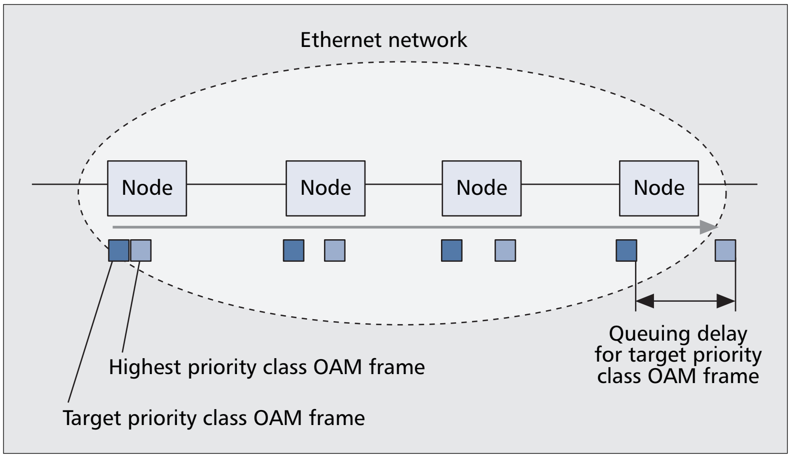

Figure 3. *High-level overview of priority-based one-way frame delay measurement scheme.*

priority class is used for network control, which requires the best buffering and scheduling treatment. The resource for the highest priority class traffic is tightly controlled and served in strict priority scheduling in each node. User traffic is transmitted in a lower priority class than network control traffic based on the quality of service (QoS) requirement of the service.

End-to-end delay consists of three parts:
- Propagation delay, D_P
- Frame processing delay, D_{FP}
- Queueing delay, D_Q

D_P is proportionally increased as the distance between source and destination increases, and D_{FP} is almost constant in a node because frame processing is implemented in high-performance hardware

The same D_{FP} and D_P values are applied to all frames that traverse the same path, regardless of the priority level of each frame. The only part that varies with priority level and network condition is D_Q. Therefore, delay performance monitoring can be done by observing D_Q values.

Figure 3 shows the concept of the priority-based one-way frame delay measurement scheme. The one-way delay measurement scheme utilizes a pair of OAM frames, which are encapsulated in two different priorities: the highest priority and the priority being monitored (target priority). They are transmitted from an MEP to a peer MEP back-to-back. The frame delay performance is determined by the interarrival time of the two frames. The traffic condition of the network and the behavior of the nodes along the path cause the delay difference between the pair pf OAM frames as they are treated differently depending on their priority levels. When the highest priority class traffic is managed to experience no congestion, the delay of the highest priority class frame is approximated as $D_P + D_{FP}$. The interarrival time between the highest priority class and the target priority class OAM frames captures the queuing delay, D_Q, of the target priority class packet. Therefore the approximate delay for each class is as follows:

Delay for the highest priority class $\cong D_P + D_{FP}$
Delay for the target priority class $= D_P + D_{FP} + D_Q$

The $D_P + D_{FP}$ value can be obtained by offline/out-of-service measurement of the round-trip delay of the highest priority class frames.

Numerical investigation shows that the inter-arrival time of the pair of OAM frames reflects the performance of the delay of the target priority class as long as we keep the delay of the highest priority class traffic constant. The measurement becomes more accurate as the traffic load increases [8, 9].

The presented scheme can measure one-way end-to-end delay accurately without the clock synchronization between two MEPs under the following two assumptions:
- The path between two MEPs is fixed.
- The network is managed so that the highest priority class traffic is not congested.

Generally speaking, a path between two nodes in data networks is not fixed, but can change over time as the network topology changes. However, the network that utlizes Y.1731 OAM is more of a carrier type of transport network, which is well managed and tightly controlled by a network operator. Even the protection scheme to achieve prompt restoration (e.g., 50 ms restoration time limit) needs both working and protection paths predetermined and fixed. As far as Y.1731 is concerned, the first assumption can be a certainty. It is not difficult to regard the second assumption as a certainty, as network operators normally assign the highest priority class to their internal network control/management task traffic. Even when the traffic volume is high due to either normal user traffic surge or malicious denial of service attacks, network operators want their internal network control traffic to get through without any delay.

Conclusion

We envision that these newly added OAM features will potentially constitute the key technologies for a carrier Ethernet transport networking because of capabilities of up to 10 Gb/s capacity with dense wavelength-division multiplexing (DWDM) system interface compatibility, SONET/SDH-level OAM and protection, WAN-level scalability, broad vertical network protocol compatibility, and broad horizontal protocol adaptability. A carrier Ethernet transport network will not only be a cost-effective data networking infrastructure, but also the most efficient technology for high performance network services with end-to-end QoS provisioning and OAM capability.

Ethernet OAM is anticipated to play an important role in defining OAM for other packet transport layer networks. The functions and mechanisms of the newly consented Y.1373 T-MPLS OAM are almost the same as those in Y.1731 Ethernet OAM. The linear protection switching for T-MPLS is very similar to Ethernet protection switching, and even their state transition tables are identical. IEEE provider backbone bridging — traffic engineering (PBB-TE) is looking to use the Ethernet OAM currently under discussion.

Recently many mechanisms have been introduced to monitor network performance and maximize network utilization. However, most of them assume every node in the network contributes to monitoring the performance of the network [10]. All the nodes along the path process and update probing packets from the ingress. This approach increases the processing load if the number of measured paths increases. On the other hand, Ethernet OAM assumes that only the ingress and egress nodes are involved in performance monitoring, and transit nodes just forward OAM frames in the same way as data frames. This path level OAM approach gains attention because it provides a unified OAM structure across circuit and packet switching.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their helpful comments and suggestions. This work was supported by the IT R&D program of MIC/IITA (2007-S-102-02, Carrier Class Ethernet Technology).

REFERENCES

[1] IEEE Draft Std. 802.1ag, "Local and Metropolitan Area Networks, Virtual Bridged Local Area Networks, Amendment 5: Connectivity Fault Management."
[2] ITU-T Rec. Y.1731, "OAM Functions and Mechanisms for Ethernet Based Networks," 2006.
[3] ITU-T Rec. G.803, "Ethernet Protection Switching," 2006.
[4] R. Aggarwal, "OAM Mechanisms in MPLS Layer 2 Transport Networks," *IEEE Commun. Mag.*, Oct. 2005.
[5] D. Mohan *et al.*, "MPLS and Ethernet OAM Interworking," draft-mohan-pwe3-mpls-eth-oam-iwk-00.txt.
[6] ITU-T Rec. Y.1373 (2007), "Operation & Maintenance Mechanism for T-MPLS Layer Networks."
[7] K. Sridhar *et al.*, "End-to-end Ethernet Connectivity Fault Management in Metro and Access Networks," White Paper, Alcatel Networks, 2005.
[8] J. Song *et al.*, "DiffProbe: One way Delay Measurement for Asynchronous Network and Control Mechanism in BcN Architecture," *Proc. 8th Int'l. Conf. Adv. Commun. Tech.*, Feb. 2006.
[9] I. Joo *et al.*, "Performance Monitoring for Multimedia Traffic using Differentiated Probe (DiffProbe)," *Proc. ICME '07*, July 2007.
[10] S. Kandula *et al.*, "Walking the Tightrope: Responsive yet Stable Traffic Engineering," *SIGCOM 2005*.

BIOGRAPHIES

JEONG-DONG RYOO (ryoo@etri.re.kr) is a principal member of research staff with the Electronics and Telecommunications Research Institute (ETRI), South Korea. He holds Master's and Ph.D. degrees in electrical engineering from Polytechnic University, Brooklyn, New York, and a Bachelor's degree in electrical engineering from Kyungpook National University, South Korea. After completing his Ph. D. study in the area of telecommunication networks and optimization, he started working for Bell Labs, Lucent Technologies, New Jersey, in 1999. While he was with Bell Labs, he was mainly involved with performance analysis/evaluation/enhancement study for various wireless and wired network systems. Since he left Bell Labs and joined ETRI in 2004, his work has been focused on next-generation network and carrier class Ethernet technology research, especially participating in OAM and protection standardization activities in ITU-T. He co-authored *TCP/IP Essentials: A Lab-Based Approach* (Cambridge University Press, 2004). He is a member of Eta Kappa Nu.

JONGTAE SONG (jsong@etri.re.kr) is a senior research staff member in the Broadband Convergence Network (BcN) research division of ETRI. He received his B.S. degree in electronics and electrical engineering from the Korea Advanced Institute of Science and Technology (KAIST) in 1990, his M.S. degree in electrical engineering from the University of Southern California in 1994, and his Ph.D. degree in electrical engineering from Polytechnic University, Brooklyn, in 1998. He worked for Bell Labs Lucent Technologies from 1998 to 2001 and for several startup companies (including Coree Networks and Parama Networks) in New Jersey from 2001 to 2004. Since he joined ETRI in 2004, his work has been focused on BcN network architecture, QoS control, and flow-based network control architecture.

JAEWOO PARK (parkjw@etri.re.kr) is a senior member of engineering staff at ETRI. He received his B.S.E.E. from Inha University, South Korea, in 1993 and his M.S.E.E. from KAIST in 1995. He specializes in network system design using network processors. His area of interest is high speed wired and wireless network architecture, especially WiMAX technology.

BHEOM-SOON JOO (bsjoo@etri.re.kr) received his B.S. degree from Seoul National University, Korea, and his M.S. degree from KAIST in electrical and electronic engineering. Since December 1983 he has been working with ETRI, participating in system development projects such as TDX10, ATM exchange, 10-gigabit Ethernet switch systems, and Carrier-class Ethernet systems. Currently, he works as a leader in the Carrier Ethernet research team. His research interests include Ethernet technology, broadband convergence networks, high-speed networks, and network synchronization.

Carrier Ethernet transport network will be not only a cost-effective data networking infrastructure but also the most efficient technology for high performance network services with end-to-end QoS provisioning and OAM capability.