Carrier ethernet OAM: An overview and comparison to IP OAM

Original

Carrier ethernet OAM: An overview and comparison to IP OAM / Hofstede, Rick; Drago, Idilio; Moura, Giovane C. M.; Pras, Aiko. - STAMPA. - 6734:(2011), pp. 112-123. (Intervento presentato al convegno 5th International Conference on Autonomous Infrastructure, Management, and Security, AIMS 2011 tenutosi a Nancy, France nel 2011) [10.1007/978-3-642-21484-4_14].

Availability:

This version is available at: 11583/2659194 since: 2019-05-06T14:54:24Z

Publisher:

Springer

Published

DOI:10.1007/978-3-642-21484-4_14

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Carrier Ethernet OAM: An Overview and Comparison to IP OAM

Rick Hofstede, Idilio Drago, Giovane C.M. Moura, Aiko Pras

University of Twente, The Netherlands

`r.j.hofstede@student.utwente.nl`, `{i.drago,g.c.m.moura,a.pras}@utwente.nl`

Abstract. Ethernet has evolved from a local area to a wide area network technology. When it is used in a service provider environment, it has more complex requirements, which demand a set of management techniques for the Ethernet layer. Ethernet OAM comprises a set of management techniques for Carrier (or Metropolitan) Ethernet networks. Carrier Ethernet devices often have IP connectivity for management purposes, which might be used for IP OAM as an alternative management solution to Ethernet OAM. This paper provides an overview of Carrier Ethernet technology and evaluates whether, and until which extent, IP-based protocols can replace Ethernet OAM in Carrier Ethernet networks.

Keywords: Network Management, Carrier Ethernet, Metropolitan Ethernet, Provider Backbone Bridging, Ethernet OAM, IP OAM

1 Introduction

Ethernet has been the number one link-layer technology in local area networks (LANs) for a long time [1]. To make Ethernet suitable to be used in larger networks, such as metropolitan and wide area networks (MANs and WANs respectively), it was extended to provide high availability, quality of service, secure communication, and superior scalability [1]. Carrier Ethernet¹ allows service providers to offer connectivity at the Ethernet level, in contrast to the IP level. Some backbone network operators are already considering a deployment of Ethernet in their Next-Generation Network [2].

With the deployment of Carrier Ethernet services it is required to manage Ethernet in an end-to-end manner, besides managing it on the link-layer. To achieve this, IEEE and ITU-T have standardized two Operations, Administration & Maintenance (OAM) protocol suites by means of IEEE 802.1ag [3] and ITU-T T.1731 [4], respectively. Ethernet OAM offers end-to-end fault and performance management over Ethernet Virtual Connections (EVCs), which are logical connections between various customer sites [5].

From a functional point of view, the Ethernet OAM standards seem to be very similar to IP OAM protocols, such as Ping and Traceroute. Assuming that

¹When Ethernet technology is used in large-scale (*e.g.* service provider) networks, it is commonly referred to as ‘Carrier Ethernet’ or ‘Metropolitan Ethernet’.

most Carrier Ethernet devices support IP OAM, one could wonder whether OAM functionality is now duplicated at both the link and network layer. Moreover, if this is the case, could IP OAM functionality be used as a replacement for Ethernet OAM? Since IP OAM has been used already for a long time, operators have more experience with it than with Ethernet OAM. This lead us to the following research questions:

1. *What exactly is Carrier Ethernet and which functionality does it provide?*
2. *How does Ethernet OAM functionality compare to IP OAM, and, more specifically, can IP-based protocols in Carrier Ethernet networks provide the same functionality as comparable Ethernet OAM management techniques?*

Since many papers have already been published on the topic of Ethernet OAM, we will first review those that focus on using alternative solutions for Ethernet OAM in Section 2. Section 3 describes the evolution of Ethernet towards a WAN protocol, especially for deployment in service provider environments. The Ethernet OAM standards for managing Carrier Ethernet networks are discussed in Section 4. Although several works state that IP OAM is not (entirely) suited to manage a Carrier Ethernet [1] [5], a clear explanation or motivation is not provided. Section 5 will fill this gap by defining an Ethernet OAM deployment scenario, by means of which an IP-based approach will be analyzed. Finally, we close this paper in Section 6, where we draw our conclusions and future work.

2 Related Work

Many studies have already been performed on the usage of Ethernet OAM for WANs. McFarland *et al.* [1] state that enterprise networks typically have straightforward topologies and that IP-based protocols such as SNMP, ICMP, Ping and Traceroute will suffice for management. However, it will not be suitable for managing service provider networks, carrying thousands of services for different customers. Motivations for the unsuitability of IP are not given.

Indukuri goes beyond IP-based protocols for managing Ethernet networks [6], by outlining the use of IP Ping, MPLS LSP Ping, Bidirectional Forwarding Detection (BFD) and especially Ethernet OAM for Virtual Circuit Connectivity Verification (VCCV). It is desirable for metropolitan and especially critical networks to have a fast and accurate fault detection mechanism. Such a sub-50 ms detection and restoration facility is provided by BFD and Ethernet OAM. The author concludes, however, that the choice for a VCCV mechanism should not only depend on technical decisions, but also on the underlying transport infrastructure. In the case that a virtual circuit is constructed on top of an Ethernet network (*e.g.* an Ethernet Virtual Connection (EVC)) and end-to-end management should be performed, it is wise to use Ethernet OAM in order to avoid the need for translation layers between different network layers. However, this work does not outline why IP-based management techniques do not suffice for the management of Carrier Ethernet networks.

3 Carrier Ethernet Evolution

During its evolution from a LAN technology to a MAN and WAN technology, Ethernet was extended to support customer traffic separation, quality of service (QoS) and, most importantly, a greater number of MAC addresses (of customers, among others) in the forwarding tables of switches [7]. The frames of the various Ethernet standards are depicted in Figure 1. ‘Ethertype’ and ‘Frame Check Sequence’ fields are left out for the sake of space. The evolution of Ethernet has been standardized by the IEEE in several standards, starting with IEEE 802.1Q [8]. This standard adds a VLAN tag to an Ethernet frame, right after the source and destination MAC addresses, by means of which the forwarding plane can be partitioned into logical segments.

In the same year as IEEE 802.1Q was standardized, an amendment was defined in IEEE 802.1ad [9], also known as Provider Bridging (PB). We assume this standard and all following ones to be Carrier Ethernet standards. The IEEE 802.1Q VLAN tag contains VLAN IDs of 12 bits, supporting up to 4094 VLANs. Although this number of VLANs will be enough for most LANs, it will not suffice for large service provider environments. To overcome this scalability problem [7], IEEE 802.1ad defines VLAN tag stacking, allowing service providers to insert an additional VLAN tag of 12 bits in an already tagged frame. This ‘S-VID’ VLAN tag is only used inside the service provider domain and is inserted in front of the initial VLAN tag, which is now referred to ‘Customer VID’ (C-VID).

IEEE 802.1ah [10], also known as Provider Backbone Bridging (PBB), allows a strict separation between customer and service provider domains by encapsulating customer traffic. This is achieved by inserting a new Ethernet header in front of the existing one, including a new backbone VLAN tag (B-VID) and a new ‘Service Instance ID’ (I-SID) field. The latter can be considered an extended VLAN ID, used to identify customer instances inside the operator network. By considering the entire PB frame as payload and inserting a new Ethernet header in front of it, a completely isolated address space is used inside the Ethernet backbone network. The result is a drastically reduced complexity/size of the forwarding tables in backbone nodes, since only backbone node addresses and backbone VIDs are needed for switching.

Fig. 1. Evolution of Ethernet frames

The most recent Carrier Ethernet standard is IEEE 802.1Qay [11]. It uses the same frames as IEEE 801.ah (PBB), but it adds traffic engineering capabilities and related rapid protection against failures [12]. IEEE 802.1Qay is therefore referred to as Provider Backbone Bridging with Traffic Engineering (PBB-TE). It adds support for static, traffic-engineered paths by replacing the use of the spanning tree protocol (STP) by an external method. Besides disabling STP, also broadcasting and MAC address learning are disabled [13]. Broadcast traffic and traffic for unknown destinations are discarded by the edge nodes of the network.

In the next section we present Ethernet OAM, by means of which fault and performance management have been added to Ethernet.

4 Ethernet OAM

End-to-end OAM has been added to Ethernet by means of IEEE 802.1ag & ITU-T Y.1731. While these standards had different focal areas when work on them started, IEEE 802.1ag is nowadays considered a subset of ITU-T Y.1731. Both standards cover fault management, while performance management is solely covered by ITU-T Y.1731. Fault management can be used for detecting and isolating faults in a network, just as notifying about faults. Performance management allows to measure throughput, delay, etc. This will help to verify and prove service performance against a Service-Level Agreement (SLA) [14], for instance.

Before the Ethernet OAM management techniques can be discussed, some terminology needs to be described. *Maintenance End-Points* (MEPs) are actively managed components, which are positioned at *Maintenance Domain* (MD) boundaries. Interconnected MEPs are called a *Maintenance Entity* (ME). A *Maintenance Entity Group* (MEG)² can include several MEs, depending on the topology: for point-to-point Ethernet connections, a MEG contains a single ME. In a multipoint setup, a MEG consists of several MEs. Inside a ME, and thus between MEPs, one or more *Maintenance Intermediate-Points* (MIPs) can be placed. MIPs only react to OAM flows, while MEPs initiate and terminate them. In order to have management hierarchies, OAM levels can be defined to run OAM mechanisms completely separated. These concepts will be highlighted again in Section 5.1, where a deployment scenario for Ethernet OAM will be discussed.

The remainder of this section will focus on the most commonly known/used Ethernet OAM management techniques.

4.1 Continuity Check

Continuity Check (CC) can be used to detect interruptions in connectivity (and thus continuity) between end points (MEPs) in an Ethernet network. This is accomplished by transmitting ‘heart-beat’ messages between MEPs, which are forwarded by MIPs. By doing so in a periodic manner, connectivity can be verified. MEPs exchange CC messages with the other MEP inside the same ME, and at each administrative OAM level.

²This terminology is based on ITU-T Y.1731. In IEEE 802.1ag, a MEG is called a Maintenance Association (MA).

4.2 Loopback

Loopback (LB) provides a way to transmit request/response messages, in order to verify bi-directional connectivity with another MEP or MIP. Upon reception of a LB message, a response message is returned towards the requester. In contrast to Continuity Check, which sends messages in a periodic fashion, LB messages are typically initiated by operator command, although nodes can be configured to transmit LB messages in a periodic fashion as well.

4.3 Link Trace

Link Trace (LT) can be used to isolate faults in Ethernet networks. MEPs send out LT messages on a particular ME, in order to identify the connectivity and relationships with remote MEPs and MIPs. While a LT can only be initiated by a MEP, all MIPs and MEPs downstream the path towards a destination MEP at the same OAM level will respond to it.

4.4 Alarm Indication Signal

Alarm Indication Signal (AIS) provides a method for notifying operators about a network anomaly. As soon as a MIP detects a failure at its OAM level, it will send out an AIS message towards the reachable peer MEPs of the same ME. After the MEPs receive the AIS from the MIPs, they will send out a multicast AIS message in the upstream direction of a fault, at the next most superior OAM level and on every service provider VLAN affected by the failure. AIS is not supported by IEEE 802.1ag.

4.5 Loss Measurement

Loss Measurement (LM) offers a way for operators to determine the amount of frame loss in an Ethernet network, over an EVC for instance. More precisely, it is the ratio between undelivered OAM frames and the total number of OAM frames transmitted during a specific time interval.

ITU-T Y.1731 defines two types of LM:

1. *Single-Ended.* LM messages are transmitted to another MEP, which includes transmission and reception frame counts in its response message. In this case, only the LM initiator is able to derive frame loss from the counters (since it does not include its local counters in the initial LM message).
2. *Dual-Ended.* Continuity Check messages are used to carry frame transmission and reception counters. In contrast to the single-ended approach, this approach allows all MEPs inside a ME to derive frame loss, instead of only the initiating node.

4.6 Delay Measurement

Delay Measurement (DM) can be used for measuring delay in a Carrier Ethernet network. The unit of measurement is the round trip delay of a frame, measured from its first transmitted bit, until the reception of its last bit. Since a DM frame needs to be sent back to its originating node, LB messages are used.

Two types of DM can be identified:

1. *One-way measurement.* An initiating MEP includes a transmission timestamp in the Ethernet frame. The destination node will capture the frame reception timestamp, and compare both timestamps. As a consequence, the clocks of the sending and receiving nodes need to be synchronized.
2. *Two-way measurement.* In contrast to the one-way measurement, this DM type does not require clock synchronization. The initiating node still includes a timestamp in the Ethernet frame. After the destination node performs a loopback on the frame, the initiating node will receive the frame again. On reception, this node will capture the reception timestamp. Finally, the difference between the timestamps can be calculated.

5 Comparison to an IP-Based Approach

The Ethernet OAM management techniques discussed in the previous section are, from a functional point of view, very similar to IP-based management techniques. A few examples:

- *Ethernet OAM Continuity Check* resembles a uni-directional IP Ping.
- *Ethernet OAM Loopback* verifies connectivity with a MEP or MIP, by performing a loopback on a frame. This is similar to IP Ping.
- *Ethernet OAM Link Trace* offers comparable functionality as IP Traceroute. Both techniques allow to trace a path between nodes through a network. Instead of using a ‘time-to-live’ (TTL) field in a frame header, MEPs/MIPs pass LT messages downstream the path towards a destination node.
- *Ethernet OAM Alarm Indication Signal* is able to send out notifications to the reachable MEPs of Maintenance Entity (ME). SNMP allows the transmission of traps/notifications to a SNMP manager as well.
- *Ethernet OAM Loss Measurement & Delay Measurement* offer similar features as certain SNMP Management Information Bases (MIBs) do.

As we outlined in the Introduction, we would like to verify whether it is possible at all, or until which extent, to manage Carrier Ethernet networks with IP-based protocols. In order to do so, IP needs to be supported on top of the Ethernet infrastructure. Although this seems to contrast the principle of having a pure Ethernet network, most network devices already have an IP interface for management purposes, in order to support a Web interface, Telnet, SSH, syslog or SNMP, for instance. Since it might not be desirable to have a full IP infrastructure on top of an Ethernet backbone network, we assume that IP will not be used for routing purposes and that all managed devices take part in a management VLAN. This has the following consequences:

1. *Only devices inside the same management VLAN and IP domain are reachable.* However, it can be advantageous for end-to-end EVC management to have nodes reachable from outside a specific domain, for instance.
2. *The IP TTL field value is not lowered on Ethernet network hop transition.*
3. *Since frames destined to a node (identified by a B-MAC and B-VID) for which no path has been defined will be discarded by ingress Backbone Edge Bridges (BEBs), paths from a management node to all managed devices need to be defined.*

In order to verify whether an IP-based approach could be used to manage a Carrier Ethernet network, we defined a typical deployment scenario for a Carrier Ethernet, in which Ethernet OAM could be deployed. It will be used to analyze an IP-based approach for managing a Carrier Ethernet network.

5.1 Deployment Scenario

Figure 2 shows our deployment scenario for a Carrier Ethernet network. Four customer sites are shown, belonging to two different customers (A and B). Customers can be end customers, service providers, operators, and access or aggregation networks [15]. Both customers acquired an EVC between their two sites. The result is a (virtual) one-hop connection between the customer-facing switch ports of the edge devices. The transport network considered here is based on IEEE 802.1Qay (PBB-TE).

A customer site is connected to a BEB, which consists of two components [15]:

1. *I-Component*: maps S-VIDs to I-SIDs (Instance IDs) and adds a PBB header with/without³ a B-VID.
2. *B-Component*: maps I-SIDs to B-VIDs and adds a B-VID to the PBB header (or the whole PBB header in case the I-Component has not done so).

Fig. 2. Deployment scenario of a PBB-TE network with Ethernet OAM

³Whether an I-Component inserts PBB header or not depends on vendor implementation.

I-Components are used for bridging in the customer space, based on customer MAC addresses and S-VIDs. B-Components are used for bridging in the provider domain based on B-MAC addresses and B-VIDs [15]. The I-Component is often called Customer Premises Equipment (CPE). A CPE is the last hop between a service provider network and a customer's equipment [16]. The two components of a BEB can be either in one or in two devices. In Figure 2, the components of the left BEB are in separate devices. EVCs are established between two CPEs.

The next hop after the first BEB is one of the Backbone Core Bridges (BCB), depending on the B-VID onto which the frame's I-SID (Instance ID) got mapped. Although a PBB network is considered here, the BCBs consider the frames as normal (VLAN-tagged) Ethernet frames. This is because the first three fields of the Ethernet header are the same for PB, PBB and IEEE 802.1Q, as shown in Figure 1. After the BCBs, the (BEB) egress switch of the backbone network is the next hop. The I-Component and B-Component are now packed into one device and perform the same tasks as the first BEB, but in reverse order.

Section 4 discussed how Carrier Ethernet networks could be managed by using either IEEE 802.1ag or ITU-T Y.1731. These standards require managed nodes to be either a Maintenance End-Point (MEP) or Maintenance Intermediate-Point (MIP). In our deployment scenario, (B-Components of the) BEBs are assigned the role of MEP and BCBs the role of MIP. Three Maintenance Entities (MEs) can be identified, one for each path through the network, so that each MEP is taking part in three MEs. In this work we consider only the operator's OAM mechanisms and a single OAM level.

The remainder of this section discusses the use of IP-based protocols in place of Ethernet OAM management techniques.

5.2 Continuity Check & Loopback

To automate Ping message transmission and to make it easier for an operator to handle this, a 'Remote Ping' MIB, which is part of the DISMAN (short for 'Distributed Management') framework, could be used. It lets a 'Local host' command a 'Remote host' to perform a Ping to a 'Target host'. Assuming that all nodes are reachable by the 'Remote host' and have a 'Remote Ping' implementation, a network operator could issue a Ping request from and to each node inside the same VLAN.

Compared to Ethernet OAM, the following advantages can be identified when using IP Ping:

- *When using IP Ping, both source and destination nodes can detect a failure.* Request messages can be sent in a periodic fashion, and a destination node must be configured to expect them in that fashion as well. When a reception timeout occurs, a faulty link or device can be assumed. This is not possible with Ethernet OAM CC, which is uni-directional by definition.
- *All nodes between which a path exists in a PBB-TE network, can exchange IP Ping messages by means of the DISMAN framework.* If such a path does not exist, frames will be dropped by the edge switches. The same is done for

broadcast traffic [13]. Ethernet OAM CC and LB can only be initiated and terminated by MEPs. Therefore, the set of nodes reachable by IP Ping can be larger than the amount of nodes reachable by Ethernet OAM.

Besides these advantages, also several disadvantages can be identified:

- *Without network-layer routing, the set of manageable nodes is restricted to a single IP/management domain.* In contrast, Ethernet OAM can be performed in an end-to-end manner over an EVC for monitoring a single service, spanning multiple domains. It is not possible for a service provider to inject (IP) packets into an EVC to verify its functioning in an end-to-end manner.
- *It is hard to ensure that IP Ping takes the path of a particular EVC, to ensure its connectivity.* When customer data arrives at an ingress BEB, it is mapped onto a B-VID, which takes a predefined path through the network. It is hard to ensure that the management VLAN uses exactly the same path.
- *A per-customer/EVC granularity requires several translation steps.* Since Ethernet OAM allows the verification of an EVC, it is immediately clear which customers are affected by a fault. More knowledge about the network is required with IP-based protocols, in order to derive the same information. Operators will need to know onto which VLAN the customer traffic is mapped. This involves active cooperation with other parties, since operators normally do not know how others mapped customer traffic onto the VLANs.

Customers can also be involved in monitoring an EVC by using IP Ping. They will, however, see the EVC service as a one-hop path. As such, they will be able to detect a problem on the EVC, but without being able to isolate it.

The deployment scenario discussed before can also be made more complex, by considering multiple operator networks between the customer sites. Customers will still see the provider domain as a one-hop connection. As soon as a customer detects a problem without Ethernet OAM, it is up to the service provider⁴ to find out which operator network causes the fault. Assuming that the path of the operator’s management VLAN is the same as the customer’s EVC, fault isolation with IP Ping is possible. With Ethernet OAM however, several (external) nodes could be configured as MIPs, so that a service provider could isolate a fault directly, even if it is located in another administrative domain.

5.3 Link Trace

In a similar way as we discussed the replacement of Continuity Check and Loopback by IP Ping, we assume the use of IP Traceroute as a replacement for Link Trace (LT). By adjusting the value in the ‘time-to-live’ (TTL) field of the IP header, the path through the network can be traced. Since IP is only used inside the management VLAN and not used for routing purposes, intermediate nodes towards a destination will never modify the TTL field value. All traces will then consist of one-hop connections and link tracing by using IP Traceroute will therefore never work in our deployment scenario.

⁴We assume the service provider network here to consist of at least two operator networks.

5.4 Alarm Indication Signal

Ethernet OAM allows the transmission of fault notifications by means of an Alarm Indication Signal (AIS). In IP-based networks, SNMP traps can be used for the transmission of notifications from agents to managers. Several default traps have been defined, such as ‘linkup’ and ‘linkdown’. When a node inside the network detects a failure, it could send out a trap to an SNMP manager.

SNMP as a replacement for AIS offers several advantages:

- *SNMP traps can be sent out to an arbitrary set of SNMP managers inside the management VLAN, while AIS can only be sent out to MEPs.*
- *SNMP offers more flexibility in defining trap structures, by allowing the definition of custom (‘enterprise-specific’) traps.* An arbitrary set of variables can be included in a SNMP trap. Also different traps can be sent for different purposes, while Ethernet OAM AIS has a fixed structure.

Besides these advantages, several disadvantages can be identified. At first, AIS can be multicasted on each S-VLAN affected by the failure automatically. This is also possible with IP-based solutions, but this requires more overhead in deriving the EVCs/customers affected by a failure. Second, SNMP traps can only be sent out inside a single IP/management domain.

5.5 Loss Measurement

Ethernet OAM Loss Measurement (LM) calculates the frame loss between two MEPs, by comparing the difference between OAM frame transmission and reception counters at the MEPs of a particular ME. By means of the RMON-MIB [17], SNMP manages Ethernet interface counters, such as ‘etherStatsPkts’. Although this counter keeps track of the sum of ingoing and outgoing frames, it is possible to define an ‘enterprise-specific’ MIB which manages these counters individually. Some MIBs exist for this purpose, such as a ‘Round Trip Time Monitoring’ (RTTMON) MIB and ‘Service Assurance Agent’ of Cisco.

The use of SNMP for LM offers several advantages:

- *Arbitrary values can be retrieved, depending on the used MIB.* The IF-MIB and RMON-MIB offer a rich set of counters and other interface statistics. Ethernet OAM LM only allows OAM frame counters to be retrieved.
- *SNMP PDUs can be sent to an arbitrary set of SNMP managers inside the management VLAN.* If all nodes have the IF-MIB/RMON-MIB deployed, a SNMP manager can retrieve the counter values from each of these nodes. Ethernet OAM LM only allows MEPs to calculate loss on a path.

Ethernet OAM LM measures OAM frame loss between MEPs inside a single ME. As such, frames coming from nodes outside the ME are not considered. To do the same with SNMP, an ‘enterprise-specific’ MIB would be needed to differentiate between frame sources or types for measuring frame loss between two network end points. Besides that, SNMP PDUs can only be transmitted inside a single management/IP domain. This has been discussed in Section 5.2.

5.6 Delay Measurement

The use of IP Ping for managing Ethernet networks has been discussed before. This protocol provides round trip delay measurements together with its results. Although Ethernet OAM DM offers some sophisticated ways to compensate for processing times at end nodes, round trip delays can be measured by using IP Ping inside the management VLAN as well. This results in the same advantages and disadvantages as described before in Section 5.2. Besides that, several SNMP MIBs have been defined for the purpose of delay measurement, such as Cisco's RTTMON MIB, as discussed in the previous subsection.

6 Conclusions

This paper presented an overview of the various Carrier Ethernet standards and the related Ethernet OAM mechanisms. By considering a specific deployment scenario for a Carrier Ethernet operator network, an IP-based approach for managing these networks has been analyzed.

In the first section of this paper, two research questions were addressed:

1. *What exactly is Carrier Ethernet and which functionality does it provide?*
Compared to the initial Ethernet standard for LANs, especially scalability improvements have been added to Ethernet. This allowed Ethernet to deal better with a greater number of MAC address in wide area networks. Due to this network scale increase, management of Ethernet became much more important than before. As a result, a set of management techniques was defined, in order to manage Ethernet on the Ethernet layer. When Ethernet is used in large-scale networks and manageable by using Ethernet OAM, it is commonly referred to as 'Carrier Ethernet' or 'Metropolitan Ethernet'.
2. *How does Ethernet OAM functionality compare to IP OAM, and, more specifically, can IP-based protocols in Carrier Ethernet networks provide the same functionality as comparable Ethernet OAM management techniques?*
From a functional point of view, the various Ethernet OAM management techniques appeared to be very similar to IP OAM protocols, such as Ping, Traceroute and SNMP. For a single operator domain, most IP-based protocols discussed in this paper are able to provide similar functionality, as their Ethernet OAM 'counterparts'. IP Traceroute is the only protocol that turned out not to be functional at all. Besides that, the scope of an IP-based approach is limited to a single management/IP domain, since network layer routing was not considered in our deployment scenario. Consequently, IP-based protocols are not deployable in an end-end-end fashion over an Ethernet Virtual Connection, which makes it impossible to verify the end-to-end service offered to a customer. IP-based protocols are therefore not a suitable replacement for Ethernet OAM in operator environments.

As future work, a multi-domain Carrier Ethernet deployment could be investigated. Besides that, the use of other OAM techniques for Carrier Ethernet networks, such as MPLS OAM or SDH OAM, could be investigated.

Acknowledgements

This research work has been supported by SURFnet's GigaPort3 project for Next-Generation Networks and the EU FP7-257513 UniverSelf Collaborative Project. Special thanks to Mark Prins from TNO-ICT for his valuable contribution to the research.

References

1. McFarland, M., Salam, S., Checker, R.: Ethernet OAM: Key Enabler for Carrier Class Metro Ethernet Services. In: IEEE Communications Magazine. Volume 43, issue 11., IEEE (November 2005) 152–157
2. SURFnet: GigaPort3 and SURFnet7. <http://www.surfnet.nl/en/innovatie/giga-port3/Pages/Default.aspx> (April 2011)
3. IEEE: 802.1ag: Connectivity Fault Management (December 2007)
4. ITU-T: Y.1731: OAM functions and mechanisms for Ethernet based networks (February 2008)
5. Chiruvolu, G., Ge, A., Elie-Dit-Cosaque, D., Ali, M., Rouyer, J.: Issues and Approaches on Extending Ethernet Beyond LANs. In: IEEE Communications Magazine. Volume 42, issue 3., IEEE (March 2004) 80–86
6. Indukuri, N.: Pseudowire VCCV - BFD vs Ethernet OAM. In: 2nd International Symposium on Advanced Networks and Telecommunication Systems. (2008) 1–3
7. Allan, D., Bragg, N., McGuire, A., Reid, A.: Ethernet as Carrier Transport Infrastructure. In: IEEE Communications Magazine. Volume 44, issue 2., IEEE (February 2006) 95–101
8. IEEE: 802.1Q: Virtual Bridged Local Area Networks (May 2006)
9. IEEE: 802.1ad: Virtual Bridged Local Area Networks - Amendment 4: Provider Bridges (May 2006)
10. IEEE: 802.1ah: Virtual Bridged Local Area Networks - Amendment 7: Provider Backbone Bridges (August 2008)
11. IEEE: 802.1Qay: Virtual Bridged Local Area Networks - Amendment 10: Provider Backbone Bridge Traffic Engineering (August 2009)
12. Bottorff, P., Saltsidis, P.: Scaling Provider Ethernet. In: IEEE Communications Magazine. Volume 46, issue 9., IEEE (September 2008) 104–109
13. Lee, W., Kim, D., Song, H.: Autonomous Client Discovery in Backbone Edge Bridges for Multipoint PBB-TE Networks. In: 12th International Conference on Advanced Communication Technology (ICACT), 2010, IEEE (April 2010) 712–716
14. Ryoo, J., Song, J., Park, J., Joo, B.: OAM and Its Performance Monitoring Mechanisms for Carrier Ethernet Transport Networks. In: IEEE Communications Magazine. Volume 46, issue 3., IEEE (March 2008) 97–103
15. Luyuan, F., Zhang, R., Taylor, M.: The Evolution of Carrier Ethernet Services — Requirements and Deployment Case Studies. In: IEEE Communications Magazine. Volume 46, issue 3., IEEE (March 2008) 69–76
16. Sofia, R.: A Survey of Advanced Ethernet Forwarding Approaches. In: IEEE Communications Surveys & Tutorials. Volume 11, issue 1., IEEE (First Quarter 2009) 92–115
17. Waldbusser, S.: Remote Network Monitoring Management Information Base. RFC 2819 (Informational) (May 2000)