

Monitoring Ethernet Links in the First Mile

Roxana Stănică*, Emil Petre**

* SC Polystart Security S.R.L., Craiova, Romania (e-mail: roxana_batm@yahoo.com)

** Department of Automatic Control, University of Craiova, Romania (e-mail: epetre@automation.ucv.ro)

Abstract: Ethernet in the First Mile (EFM) specifies the protocols and Ethernet interfaces for using Ethernet over access links as a first-mile technology and transforming it into a highly reliable technology. EFM standard defines the Operation, Administration, and Maintenance (OAM) feature as an optional sub-layer just above the Ethernet Media Access Controller (MAC). This paper presents a new EFM protocol structure (for monitoring and troubleshooting the network) that is based on OAM feature. The OAM structure permits the configuration of different parameters such as timers, link connectivity, packet's priority, latency, frame-loss, and jitter. By using OAM, service providers can measure and ensure Service Level Agreements (SLA).

Keywords: Ethernet first mile, OAM, timer, ping, loopback, display

1. INTRODUCTION

Ethernet the First Mile (EFM) together with Metro Ethernet Forum (MEF) [5] are standard protocols that are based on Operation, Administration, and Maintenance (OAM) feature.

Ethernet in the First Mile (EFM), also known as IEEE 802.3ah, is a collection of protocols specified in IEEE 802.3, defining the Ethernet in the access networks, i.e. first or last mile.

Metro Ethernet Forum (MEF) standard estimates Service Assurance (SA) metrics and specific connectivity tests.

The structure of OAM feature is presented in Fig. 1 [3, 7]. OAM feature [6] describes the tools and utilities for installing, monitoring, and troubleshooting the network. This feature is a family of standards providing reliable remotely-managed Service Assurance (SA) mechanisms for both the provider and customer networks. OAM offers the ability to perform automatic periodic network-wide service assurance and quality verifications [8].

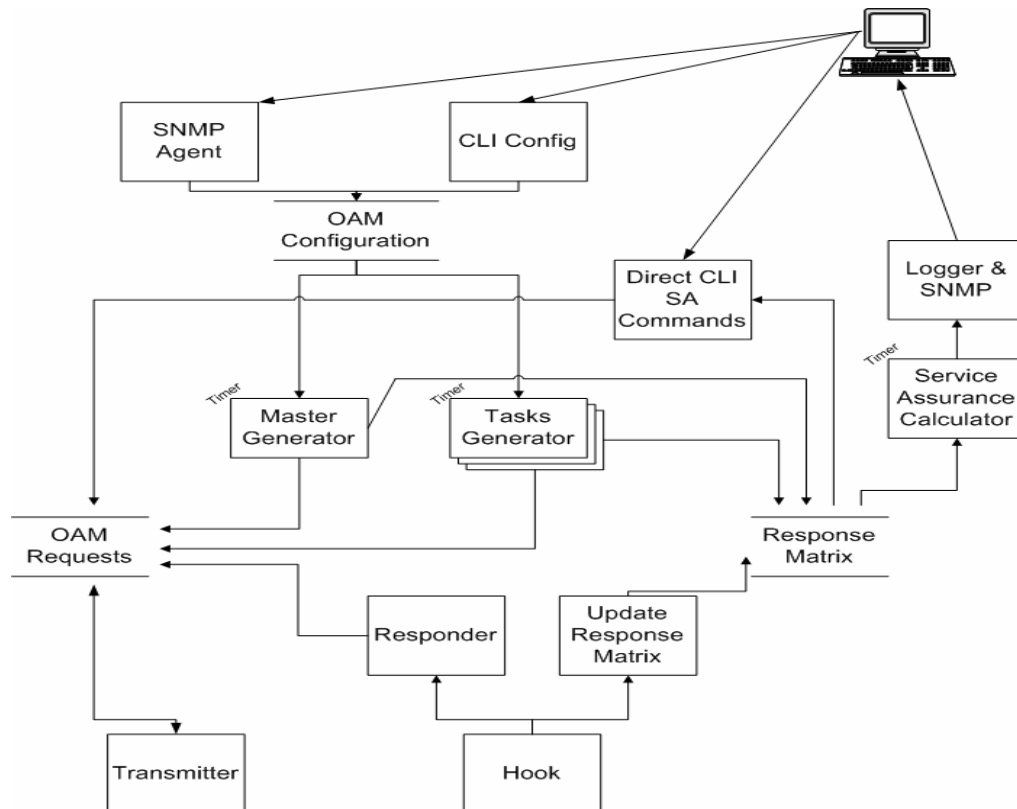


Fig. 1. OAM structure

1.1. Service Assurance (SA) Calculator Process (Timer)

This process runs once a period of time (each second for example) and each time recalculates SA statistics.

This process maintains internal store specifying latency, jitter, frame-loss or connectivity for each MEF-OAM host [5].

1.2. Master Generator (Timer)

Master Generator is responsible for all the specific OAM domain processes in the system, including coordination between task generators:

- Master Generator is responsible for the creation, destruction and updating of task generators (single SA tasks between the specific device and current device). Master generator examines the changes in OAM configuration database and destroys or creates task generators.
- Master Generator is responsible for generating periodic Hello packets [5] for the entire network. Hello packets, created by this generator, are stored in OAM Requests and Response Matrix (empty container created for each known host).
- In case of critical system/port failures, Master Generator is responsible for sending EFM-OAM event notifications.
- For all EFM-OAM Master ports, Master Generator is responsible to send EFM-OAM Protocol Data Units (PDUs) for each pre-configured period of time.

1.3. Tasks Generator (Timer)

Task Generator runs specific SA tests between current device and other specific device to determine the jitter, frame-loss, connectivity, and latency.

Connectivity Test Requests created by this generator are stored in OAM Requests and Response Matrix (empty container created in Response Matrix); see Fig. 1.

Each Task Generator runs once in a specified period, and posts its requests to the store mentioned above. If some service assurance parameter, for this task, becomes invalid (reconfigured through CLI - Command Line Interface or SNMP – Simple Network Management Protocol), this process is signalled by the Master Generator process and terminates.

1.4. Hook

This process receives all the packets with the configured EtherType and classifies them whether they need response or/and need to be stored for service assurance calculations (see Fig. 1).

- If they need response, Responder function is invoked;
- If they need to be stored, Update Response Matrix is invoked.

The following OAMPDUs need response:

- MEF-OAM Connectivity Test Request;
- EFM-OAM Information OAMPDU.

OAMPDUs are standard length Ethernet frames. They are untagged and with a length of 64 to 1518 bytes. The maximum PDU size is determined during the discovery process (see chapter II: Metro Ethernet Forum (MEF)).

EtherType is a two-octet field in an Ethernet frame. This field can also be used to represent the payload size of the Ethernet frame.

1.5. Responder

Responder is called by the hook. For ingress EFM-OAM packets (incoming packets), if the source port is in the slave mode, an EFM-OAM information packet reply is sent. For ingress MEF-OAM packets, responder performs forwarding decisions.

2. METRO ETHERNET FORUM (MEF)

Metro Ethernet Forum (MEF) protocol is described as a bridged-Ethernet per-VLAN (Virtual Local Area Network) implementation and with the help of Ping and Hello packets [5] estimates service assurance (SA) metrics and specific connectivity tests. In addition, traceroute is available for finding and displaying an Ethernet-level path from the current device to a specified device in the network.

MEF protocol is an Operations, Administration, and Maintenance (OAM) feature. MEF is a powerful tool for network engineers to test real-time connectivity, latency, jitter, and frame-loss parameters [10] on the bridged network, without higher-layer implementations.

MEF protocol provides mechanisms to perform Ethernet network-wide service assurance for providers and end-to-end service assurance for Ethernet customers. Bridges, implementing MEF-OAM standard [10], are added to the MEF-OAM domains [5] on the 802.1q bridged network and participate in service-level assurance (SLA) and fault isolation. By enabling MEF, users can monitor the activity of a single link or the entire bridged-Ethernet network.

MEF-OAM supports the following functions [5, 7]:

- Network discovery (802.1q bridged network-discovery per VLAN);
- Connectivity verification;
- Latency and loss measurement;
- Delay variation measurement;
- Traceroute functions.

3. ETHERNET in the FIRST MILE (EFM)

The IEEE 802.3ah Ethernet in the First Mile (EFM) standard specifies the Ethernet interfaces for using Ethernet over access links as a first-mile technology and transforming it into a highly reliable technology.

By using standard EFM [1], broadcast Internet access is gained in addition to services (such as Layer 2 transparent

LAN services, Voice services over Ethernet Access networks, Video, and multicast applications) reinforced by security and Quality of Service (QoS) [4] control to build a scalable network.

The in-band management specified by this standard defines the OAM mechanism needed for the advanced monitoring and maintenance of Ethernet links in the first mile. OAM capabilities facilitate network operation and troubleshooting for both the provider and the customer networks.

Basic 802.3 packets convey OAM data between two ends of a physical link. 802.3ah (Clause 57) provides the single-link OAM capabilities [8].

When EFM-OAM protocol is enabled, two connected OAM devices exchange OAMPDUs. OAMPDUs are standard-size frames that include information for the destination MAC address, EtherType, and subtype. They are sent at a predefined rate (a limitation necessary for reducing the impact on bandwidth)

3.1. Standard EFM-OAM Structure

In Fig. 2 is presented the standard EFM-OAM structure. It consists of three tasks that access the internal database together with the CLI (Command Line Interface) task.

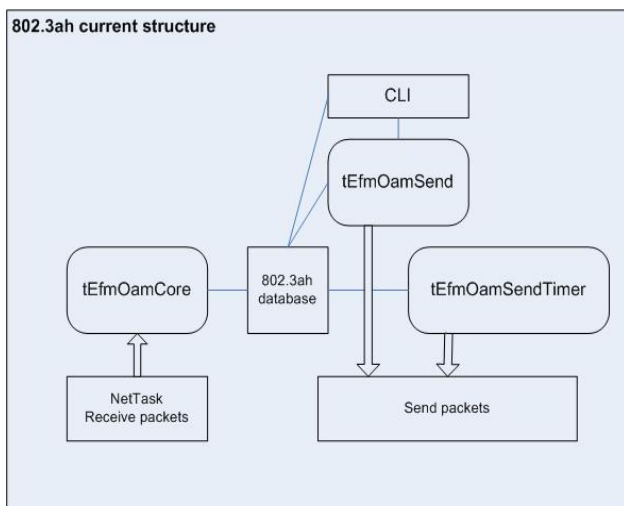


Fig. 2. Standard EFM-OAM structure

3.2. New Proposed EFM-OAM Structure

In Fig.3 is presented the new EFM-OAM. This structure has only one task tEfmOamMain (see Fig. 3) that accesses the internal database for a semaphore synchronization. The user communicates with this task using Inter Process Communication (IPC) calls. NetTask (when packets are received) sends the packets to EFM-OAM task using a communication pipe.

3.3. Connectivity timer

- Replaces the tEfmOamSendTimer (see Fig. 2) task.
- Its value is set to 100 milliseconds (minimum value of the Hello interval).
- Is started when EFM-OAM application starts (EFM-OAM is globally enabled on the device).

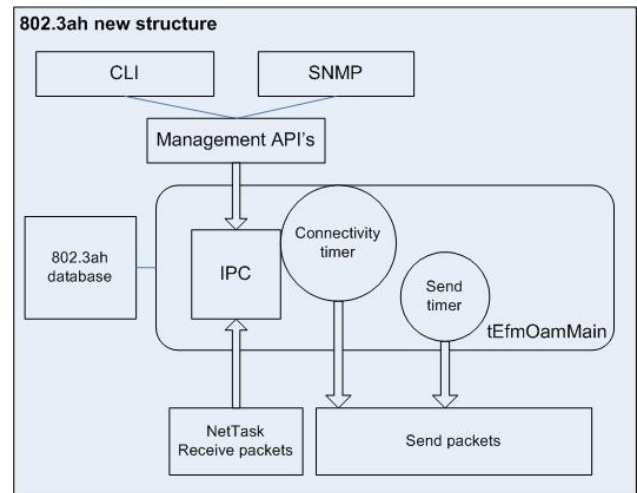


Fig. 3. New EFM-OAM structure

3.4. Send timer

- Replaces the tEfmOamSend (see Fig. 2) task.
- Assures the sending of maximum 10 packets/second.
- Its value is set to 1 second.
- Is started when EFM-OAM application starts (EFM-OAM is globally enabled on the device).

3.5. Potential Applications

By using the Ethernet demarcation service, providers can manage remote devices without using an IP layer. Instead, they can use link-layer SNMP counters request and reply, loopback testing, and other techniques that are controlled remotely.

3.6. EFM-OAM Features

EFM-OAM supports the following features [7]:

- Discovery: the local Data Terminating Entity's (DTE) ability is to discover other EFM-OAM, to enable DTEs, and exchange information about OAM entities, capabilities, and configuration.
- Link monitoring: detecting and indicating link faults to peers.
- Remote failure detection: conveying error conditions to peers via a flag in the OAMPDUs.
- Remote loopback: troubleshooting problematic segments by sending Loopback Control OAMPDUs to a peer.
- Organizing specific enhancements: providing vendor-specific enhancements to the protocol.

3.7. Timers

Two configurable timers control the protocol:

- The Hello timer determines the rate for sending OAMPDUs [5].
- The Keep-alive timer determines the time interval between two consecutive OAMPDUs sent from the peer.

An additional 1-second non-configurable timer can be used for error aggregation necessary for the Link Monitoring Process to generate link quality events.

3.8. *Flags*

Each OAMPDU includes a *Flags* field that includes the discovery process status. There are three possible status values [7]:

- Discovering: the discovery process is in progress.
- Stable: the discovery process is completed and the remote device can start sending any type of OAMPDU.
- Unsatisfied: appears when mismatches exist in OAM configuration and prevent OAM from completing the discovery process.

4. EFM-OAM FUNCTIONALITY

Two industry groups are working closely together to create and promote one global standard that will guarantee total interoperability: the Institute of Electrical and Electronics Engineers, Inc. (IEEE) 802.3ah Task Force and the Metro Ethernet Forum (MEF). While the 802.3ah Task Force is working on developing the EFM standard, the Metro Ethernet Forum is building market awareness, driving consensus, and preparing interoperability demonstrations [9].

The first mile (also called the last mile, the subscriber access network, or the local loop) is the link between the end user and the public network. It is the communications infrastructure of the business park, the neighborhood, or a single building (such as an office building, hotel or apartment). Essentially, the first mile is a figure of speech for the link from the subscriber or user to the network; therefore, it may be greater or less than a mile [9].

4.1. *Discovery*

EFM-OAM enabled DTEs identify other DTEs along with their OAM capabilities using Information OAMPDUs, advertising the following information:

- OAM configuration (capabilities): the local DTE's OAM capabilities. Using this information, a peer can determine what functions are supported and accessible (for example, loopback capability).
- OAM mode: the DTE's OAM mode, also used to determine the DTE's functionality:
 - Active mode: the DTE instigates OAM communications issues queries and commands to the remote device.
 - Passive mode: the DTE generally waits for the peer DTE to instigate OAM communications and responds to them. It does not instigate commands and queries.
- OAMPDU configuration, including the maximum size of OAMPDUs delivered (This information, in combination with a limited rate of ten frames per second, is used to limit the bandwidth allocated to OAM traffic)

- Platform identity is a combination of an Organization Unique Identifier (OUI, the first three bytes of the MAC address) and 32-bits of vendor-specific information. OUI allocation is controlled by the IEEE.

Once OAM support is detected and the OAM expectations are met, both ends of the link exchange the above information, enabling OAM on the link. However, the loss of a link or a failure to receive OAMPDUs for a predefined interval causes the discovery process to start again.

4.2. *Link Monitoring Process*

The Link Monitoring process is used for monitoring the link for occurrences where defined thresholds are crossed and notifying the remote device by sending Event Notification OAMPDUs.

The events the Link Monitoring process indicates:

- Errored Symbol per second, if the number of symbol errors that occurred during a specified period exceeded a threshold. These are coding symbol errors (for example, a violation of 4B/5B coding - each forth-bit of received data has an extra fifth bit added).
- Errored Frame per second, if the number of frame errors detected during a specified period exceeded a threshold.
- Errored Frame per N frames, if the number of frame errors within the last N frames exceeded a threshold.
- Errored Seconds Summary (errored seconds per M seconds), if the number of errored seconds (one second intervals with at least one frame error) per M seconds exceeded a threshold.

Since 802.3ah OAM does not guarantee the delivery of OAMPDUs, the Event Notification OAMPDU can be sent multiple times to reduce the probability of losing these notifications using a sequence number in order to recognize duplicate events.

4.3. *Remote Failure Indication*

A flag in the OAMPDU allows an OAM entity to send failure conditions to its peer. The failure conditions are defined as follows:

- Link Fault: this condition is detected when the receiver loses the signal. This condition is sent once per second in the Information OAMPDU.
- Dying Gasp: this condition is detected when the receiver goes down.
- Critical Event: when a critical event occurs, the device is unavailable as a result of malfunction, and has to be restarted manually. The critical events can be sent immediately and continually.

When a link does not receive signals from its peer at the physical layer (for example, if the peer's laser is malfunctioning), the local entity sets this flag to let the peer know that it's transmit path is inoperable.

4.4. Remote Loopback

In order to verify the quality of links, estimating whether a network segment satisfies a Service Level Agreements (SLA), and when troubleshooting, the active device can enable the remote peer's loopback mode, using Loopback Control OAMPDUs.

In a loopback mode, the peer loops back all the traffic (except for OAMPDU traffic and pause frames) without changing it. The remote peer acknowledges the loopback by responding with an Information OAMPDU, indicating the loopback status in the *State* field (for an example, see **show efm 1/1/1** command).

There are two kinds of loopback tests:

- Loopback using multiple ping packets (1 to 200 packets). This feature tests and displays also the local and remote peer's counters.
- Loopback using hardware-created frames at wire-speed, allowing the testing of the link under extreme high-load conditions. (These frames are discarded on the active device when they get back from the remote peer.) This feature tests and displays also the local and remote peer's counters.

5. NEW EFM-OAM COMMANDS

Fig. 4 details the steps required to configure the EFM parameters [7]. To apply and configure the new EFM commands, the steps displayed in Fig.4 must be followed.

5.1. Enable or disable EFM-OAM

- **efm enable** command [7] enables the EFM-OAM protocol on the devices.

- **efm disable** command [7] disables the EFM-OAM protocol.

With **efm disable/enable** command, all EFM-OAM parameters can be configured to their default values.

5.2. Monitoring, Testing, and Loopback

The following are the new proposed EFM-OAM monitoring and network testing commands:

- **efm ping UU/SS/PP [number <number>] [delay <delay>] [timeout <timeout>]** command enables the EFM-OAM non-intrusive monitoring of a specific interface.
 - *UU/SS/PP* represents the unit, slot, and port numbers of the configured interface.
 - *<number>* defines the number of echo packets to be sent, in the range of <1–10> (default of 5 packets).
 - *<delay>* defines the delay between packets, in the range of <1–600> seconds (by default there is no delay).
 - *<timeout>* defines the reply timeout in the range of <1–60> seconds (default of 2 seconds).
- **efm loopback UU/SS/PP** command enables EFM-OAM monitoring of a specific interface (*UU/SS/PP*), by setting the remote device into a loopback mode and generating test traffic.

One of the two loopback modes can be selected:

- 1). *Storm* mode sets the remote peer interface into a loopback mode, stops the local data flow to this interface, and the local CPU generates a packet burst. When the remote peer sends the burst back, the local device validates it and displays the burst statistics.

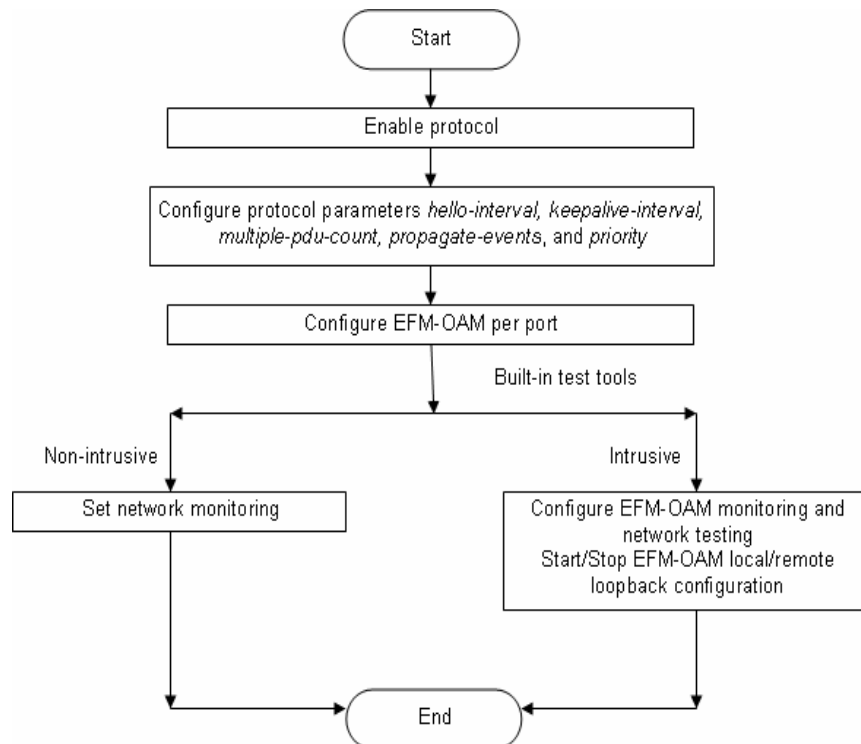


Fig. 4. EFM-OAM configuration structure

- efm loopback *UU/SS/PP storm* [delay <delay>] [packet-size <packet-size>] [no-loopback] [timeout <timeout>]
 - *storm*: selects a storm loopback.
 - <delay> defines the delay between packets in the range of <1–600> seconds (by default there is no delay).
 - <packet-size> defines the test-packets' size, in the range of <64–1512> bytes (default size of 64 bytes).
 - *no-loopback* does not define a remote loopback for this operation (set the loopback manually)
 - <timeout> defines the reply timeout, in the range of <1–600> seconds (default of 2 seconds).

2). *Burst* mode sets the remote peer interface into a loopback mode, stops the local data flow on this interface, and the local hardware generates a test packet burst (a single packet, generated by local CPU, is repetitively sent by the hardware). When the remote peer sends the burst back, the local device ignores it and displays only counters

- efm loopback *UU/SS/PP burst* [packet-size <packet-size>] [no-loopback]
 - *burst* selects a burst loopback.
 - <packet-size> defines the test-packets' size, in the range of <64–1512> bytes (default size of 64 bytes).
 - *no-loopback* does not define a remote loopback for this operation (set the loopback manually)

5.3. Clear and Display

To see all changes performed or the delete the EFM-OAM history, the following commands can be used:

- clear efm history command deletes the EFM-OAM buffer history contents.
- show efm [*UU/SS/PP*] command displays the current EFM-OAM configuration and status for a specific interface or for all interfaces.
- show efm history [*UU/SS/PP*] command displays the history for a specified interface or for all interfaces.

5.4. Examples

#show efm 1/1/1

```

Events sending status: Enabled
Event Notification Duplication Count: 5
Intervals: Keep-Alive is 5000 milliseconds, Hello is 1000 milliseconds
History limit: 24 hours or 5000 entries
Local MAC: 00:15:16:27:12:40
Efm-Oam Pkts counter : sent = 106776 ,
received = 377734
Port |Local |Remote MAC |Remote|Remote |Local
-----+-----+-----+-----+-----+-----
1/1/1 |Active |00:15:16:27:14:23|Passive|Stable |Stable
-----+-----+-----+-----+-----+-----

```

#show efm history

```

3/1/2011 19:20: Port 1/1/1: Remote Link Fault Bit Received
3/1/2011 19:21: Port 1/1/1: Remote Errored Frame Event Received
Timestamp: 12323445 Window: 30 sec
Threshold: 50 Errors: 55
Total Errors: 78654
Total Events: 9943
3/2/2011 19:21: Port 1/1/1: Remote Link Fault Bit Cleared
4/2/2011 22:30, Port 1/2/2: Remote Errored Frame Event Sent
Timestamp: 24523445 Window: 45 sec
Threshold: 10 Errors: 15
Total Errors: 32654
Total Events: 5943
3/4/2011 13:25, Port 1/1/1: Dying Gasp Received
3/4/2011 13:26, Port 1/1/1: Renegotiation Completed.
3/4/2011 13:27, Port 1/1/1: Unknown Organization
Specific Event

```

#clear efm history

#show efm history

```
%Efm-Oam history empty
```

#efm loopback 1/1/1 storm packet-size 64

```

Setting Loopback ..... Started .... Completed
Generating Test Traffic ..... Started .... Completed
Sent: 100 packets / 6400 octets
Received Successfully: 99 packets / 6336 octets

```

	Local		Remote
InOctets	636728	InOctets	1005096
OutOctets	613104	OutOctets	1136751
InUcastPkts	7500	InUcastPkts	7700
InNUcastPkts	2250	InNUcastPkts	7983
OutUcastPkts	7400		
OutNUcastPkts	2176		
InDiscards	0		
OutDiscards	0		
InErrors	0		
OutErrors	0		

#efm loopback I/I/I burst

Setting Loopback	Started	Completed
Stopping loopback	Started	Completed
Sent:	1488000 packets / 611328816 octets			
Received Successfully:	1485675 packets / 611303112 octets			
	Local		Remote	
InOctets	1669371083	InOctets	3910908339	
OutOctets	632358980	OutOctets	1669699696	
InUcastPkts	565339720	InUcastPkts	3223506341	
InNUcastPkts	26540	InNUcastPkts	1086852153	
OutUcastPkts	402271			
OutNUcastPkts	290145			
InDiscards	0			
OutDiscards	0			
InErrors	0			
OutErrors	0			

6. CONCLUDING REMARKS

High bandwidth capabilities and effective solutions [8], in terms of costs, are engines that make Ethernet a natural choice for these services. In addition, critical business applications, voice, and video conferencing require specified levels of service quality to guarantee application performance.

The operation of OAM on an Ethernet interface does not adversely affect data traffic as OAM is a slow protocol with very limited bandwidth potential, and it is not required for normal link operation. This slow protocol can be implemented in hardware or software, ensuring media independence. By using the slow protocol MAC address, OAM frames are intercepted by the MAC sub-layer and cannot propagate across multiple hops in an Ethernet network.

The first mile is the link between the end user and the public network. This link is vital for the delivery of broadband applications to residential and business subscribers [9].

Intermediate 802.3ah EFM-OAM standard specifies the protocols and Ethernet interfaces for using Ethernet over

access links as a first-mile technology and transforming it into a highly reliable technology. EFM allows carriers to extend the reach of their high performance Ethernet networks over existing infrastructure as well as new media.

In this paper, a new structure was proposed for EFM protocol based on OAM feature. New commands were added for configuring the EFM-OAM monitoring, testing and loopback processes, as well as new commands for displaying the changes performed.

7. REFERENCES

- [1] M. Beck, "Ethernet in the First Mile: The IEEE802.3ah EFM Standard," McGraw-Hill Professional Publishing, pp. 200-304, 2005.
- [2] W. Diab, M. F. Howard "Ethernet in the First Mile: Access for Everyone," IEEE Press - Standards Information Network, pp. 37-63, 2006
- [3] R. Stănică, E. Petre, "Operation, Administration, and Maintenance," Annals of the University of Craiova, vol. 6(33), no. 2, pp. 80-86, 2009
- [4] R. Stănică, E. Petre, "Control the Packets Transmission Using Quality of Service Protocol", Proceedings of 14th International Conference on System Theory and Control (Joint Conference of SINTES14, SACCS10, SIMSIS14), pp. 535-540, Sinaia, Romania, 2010
- [5] R. Stănică, E. Petre, "Control Two-Way Monitoring Process Using Metro Ethernet Forum", Annals of the University of Craiova, vol. 7(34), no. 1, pp. 66-72, 2010
- [6] M. Vigoureux, D. Ward, and M. Betts, "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks," Internet Engineering Task Force (IETF). RFC 5860. pp 3-15, 2010.
- [7] <http://www.batm.com/int/index/en/sdoc/type/111>
- [8] http://www.cisco.com/en/US/prod/collateral/routers/ps368/prod_white_paper0900aecd804a0266.html
- [9] <http://metroethernetforum.org/PDFs/EFMA/ethernet%20first%20mile%20overview%20v3.pdf>; "reproduced with permission of the Metro Ethernet Forum"
- [10] http://metroethernetforum.org/MSWord_Documents/MEF10.doc; "reproduced with permission of the Metro Ethernet Forum"