

Operation, Administration and Maintenance

Roxana Stănică*, Emil Petre**

*S.C. BATM Systems S.R.L., Craiova, Romania (e-mail: rstanica@batmsystems.ro)

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

Abstract: By using Operations, Administration and Maintenance (OAM) tools, service providers can measure and ensure Service Level Agreements (SLA). This paper presents two protocols for monitoring and troubleshooting the network, and their structure that is based on OAM tool. The OAM structure permits the configuration of different parameters such as timers, link connectivity, packet's priority, latency, frame-loss, and jitter. The standard Metro Ethernet Forum (MEF) protocol was designed to support the OAM tool and a new structure was developed for Ethernet in the First Mile (EFM) protocol.

Keywords: administration, monitoring, structure, Ethernet, connectivity, timer.

1. INTRODUCTION

Operation, Administration and Maintenance (OAM) feature 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 [6].

Two of the standard protocols that are based on OAM

feature are [10]:

- Metro Ethernet Forum (MEF) standard that estimates SA metrics and specific connectivity tests.
- Ethernet in the First Mile (EFM) standard (also called IEEE Std 802.3ah) that defines the OAM tool as an optional sublayer just above the Ethernet Media Access Controller (MAC).

The structure of OAM feature is presented in Fig. 1 [10]:

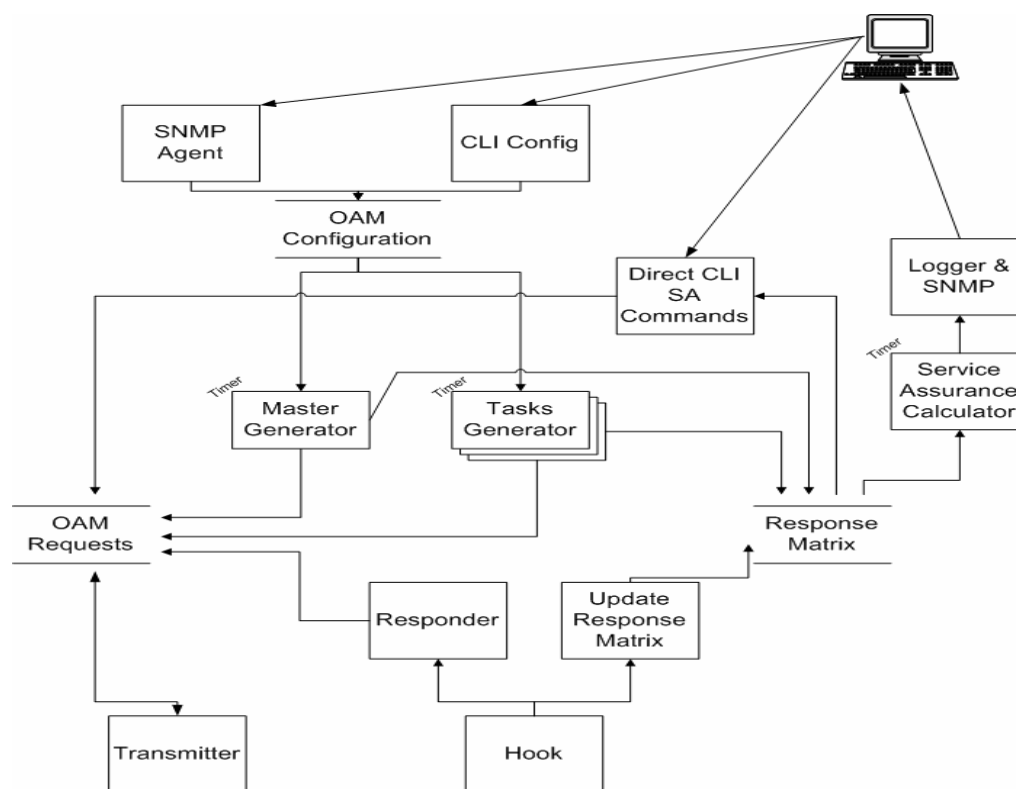


Fig. 1. OAM structure

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.

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 for 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.

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).

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), this process is signalled by the Master Generator process and terminates.

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.

- 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.

Responder

Responder is called by the hook. For ingress EFM-OAM 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)

MEF-OAM is described in [10] as implementing bridged-Ethernet per-VLAN (Virtual LAN) network-wide using Ping and Hello packets for estimating service assurance 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.

By enabling MEF-OAM, you gain the ability to monitor the activity of a single-link or the entire bridged-Ethernet network. This OAM feature is a powerful tool for testing real-time connectivity, latency, jitter, and frame-loss parameters [9] on the bridged network, without a higher-layer implementation.

MEF-OAM mechanisms perform Ethernet service assurance for providers and end-to-end service assurance for Ethernet customers. Bridges are added to the MEF-OAM domains on the 802.1q bridged network and participate in Service Level Assurance (SLA) and fault isolation.

MEF-OAM supports the following functions [8, 10]:

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

The direct effect of OAM service assurance is the discovery of the equipment on each VLAN and the maintenance of connectivity tables regarding this equipment.

MEF-OAM Domains

The MEF-OAM domain is defined by a connected topology of bridges, in which OAM operates unobstructed. At the edges of a domain, filtering constrains prevent Ethernet OAM frames from exiting the domain. On each bridge, each port is classified as interior or exterior for that domain. Up to seven domains can co-exist on a device.

MEF-OAM PDUs Generation and Forwarding Rules

MEF-OAM PDU exchange is the key element of the MEF-OAM implementation. Forwarding decisions reflect

connectivity discovery and security requirements of the implementation.

You can tag the MEF-OAM PDUs for a specific VLAN. Untagged PDUs belong to the default VLAN assigned to the port that received the PDU.

Each bridge defines a MEF-OAM domain on specific ports and VLANs.

The MEF-OAM domain on each bridge is specified by [7]:

- the VLANs on which OAMPDUs are forwarded;
- the VLAN ports participating in the OAMPDU exchange;
- the maximum allowed OAMPDU exchange rate between all MAC addresses in this domain.

Generating MEF-OAM PDUs

When a bridge generates a multicast OAMPDU on a specific MEF-OAM domain:

- The bridge must send this PDU to all allowed VLANs on all allowed ports for that domain.
- The bridge must receive PDUs up to the maximum specified rate from the ports on the same VLANs.
- If the generated packet is an OAM multicast response, the bridge sends the PDU to the port on which the destination MAC address, for a specific VLAN, is learned (using its FDB – Forwarding Data Base).

Unicast OAMPDU

When a bridge generates a unicast OAMPDU on a specific MEF-OAM domain:

- In case a VLAN is specified, the bridge must send the PDU from the port that has connectivity to the destination MAC address on that VLAN (using its FDB). If no adequate entry exists in the FDB, the bridge must send this PDU to all allowed ports on that VLAN.
- If no VLAN is specified, repeat the above step for all VLANs.

Forwarding OAMPDUs

When receiving an OAMPDU on a bridge:

- If the device is not part of the MEF-OAM domain specified in the OAMPDU, the device forwards this packet as a normal Ethernet packet.
- If the device is a part of the MEF-OAM domain specified in the OAMPDU, the below rules are followed:
 - If this is a unicast packet, the bridge checks if the VLAN specified in the OAMPDU and the target port determined by the FDB are members of the MEF-OAM domain. If they are, the bridge forwards this packet as a normal Ethernet packet. Otherwise, the bridge broadcasts this packet (unknown unicast).
 - If the received packet is an OAM multicast response, the bridge forwards the PDU to the port on which the destination MAC address for the specific VLAN is learned (using its FDB).

- If this is an OAM multicast packet but not a Connectivity Response, the bridge extracts the VLAN number from the OAMPDU and forwards it to all the ports defined by the OAM process in that specific VLAN (using Spanning Tree Protocol (STP) multicast considerations).

Network Discovery

At the Discovery stage, each device initiates periodic Hello multicasts. Each device in MEF-OAM domain receiving these multicasts builds a connectivity table and a jitter table of the reachable bridges within the domain.

Connectivity Verification

The connectivity verification process sends a Connectivity Test Request to the specified bridge and waits for a connectivity test reply message.

The alternative mechanism is to use periodic Hello frames between the devices within MEF-OAM domain. Each Hello frame between two MEF-OAM domains indicates an existing connectivity from the source bridge to the target bridge.

To measure specific service assurance aspects, the system administrator can initiate a connectivity verification test, in which one device sends a Connectivity Test Request to a remote bridge. Each device in the network that supports MEF-OAM replies to the Connectivity Test Request with a Connectivity Test Response. Based on this response, the device originating the test calculates the two-way packet loss, two-way packet jitter, and two-way packet delay.

Service Assurance Metrics Calculation

Service assurance metrics are calculated per VLAN. In case where several VLANs exist, the service assurance metrics between them is calculated several times - for each entity on the same MEF-OAM domain that shares the same VLANs.

To ensure service assurance, administrators can define threshold metrics values for different conditions (such as one-way connectivity and jitter, two-way connectivity, jitter, frame-loss, and latency) and the allowed number of consecutive samples that exceeds them. The tests results are compared to the pre-defined thresholds generating warning or error messages.

The administrators can also define the creation of log lines in case the conditions of thresholds are met, indicating that specific alarms are not triggered.

One-Way Connectivity and Jitter

Each MEF-OAM domain defined on the device can multicast periodic Hello packets (the multicast period, the packet size and their quantity are configured through the CLI -Command Line Interface).

All devices that belong to the same MEF-OAM domain and receive these PDUs can:

- forward these PDUs as multicast;

- add the PDU record to the cyclic queue (a bucket with a pre-defined size) of received OAMPDUs per MAC address and VLAN;
- periodically update the Connectivity Matrix, specifying the local timestamp of last Hello packet received for the specific device on the specific VLAN, creating the Connectivity Test Result.

Two-way Connectivity, Jitter, Frame-Loss, and Latency

By using the two-way connectivity, administrators can test and monitor the connectivity parameters between a device and a specific MAC address, or all members of the MEF-OAM domain on a specific VLAN.

This test is not issued by default and has to be triggered by configuring the following test parameters:

- the MEF-OAM domain to which this process belongs;
- the destination MAC address or multicast (all MEF-OAM domain members of the specific VLAN);
- the test period (for example, every 60 seconds);
- the size and number of packets sent per each test.

When sending a Connectivity Test Request, the source device stores a copy of this request in the response bucket. This copy exists for timeout validation in case no response arrives.

For spreading the Connectivity Test Request multicast, administrators can manually execute OAM ping.

All the devices receiving a Connectivity Test Request:

- forward the PDU if it is a multicast;
- stop processing the PDU if it is not directed to the current device;
- if the PDU is directed to the current device, send a unicast Connectivity Test Response OAMPDU to the source device;
- update the PDU timestamp for this device on the specific VLAN;

When the source device receives a Connectivity Test Response, the test:

- updates the timestamp of this device;
- adds a record of this PDU, for the specific MAC address on the specific VLAN, to the connectivity cyclic queue of the received OAMPDUs;
- calculates the jitter for this MAC address on the specific VLAN (contained in this PDU), comparing its timestamp to the local received timestamps of previous connectivity test responses (the variance between the packet's sending timestamp and receiving timestamp);
- calculates the latency per bucket (the differences between sent connectivity tests timestamp and connectivity tests local received timestamps);
- calculates the frame-loss per bucket (the differences between connectivity tests sent and connectivity test responses received).

In case of insufficient network performance, SA metrics can trigger the generation of log entries on a bridge or transfer them to a Syslog server.

MEF-OAM Error Messages

In order to maintain and monitor the network, MEF-OAM generates various log entries notifying of communication failure, exceeded thresholds and MEF-OAM PDU process errors.

The following error message is sent when MEF-OAM PDU is received from un-allowed port or if VLAN is not configured first:

Error: Un-allowed PDU received for domain <domain> on port UU/SS/PP on VLAN <vlan-id>

The following error message is sent when an invalid MEF-OAM PDU is received

Error: Invalid MEF-OAM packet received: <pdu-type>

The following error message is sent when more PDUs are received per second than a limit defined on the device. To remove this error, check the PDU limitation settings for the MEF-OAM domain:

Error: PDU Overflow for domain <domain>. Cause: Low PDU limit or Security attack

MEF-OAM Purpose

The most widely used MEF implementation is for CE-to-CE (Customer Edge) or PE-to-PE (Provider Edge) service assurance [11]. Below are the existing service assurance possibilities:

- The network administrator defines general thresholds for the network behaviour (network-wide connectivity, network-wide jitter).
- The network administrator defines specific thresholds between pairs of specific nodes (two-way frame-loss, delay, and jitter). Each specific threshold measurement between two MAC nodes is called a Monitoring process.
- The network administrator sends a specific Connectivity Test Request (CLI Ethernet connectivity test).

MEF-OAM Structure

The MEF-OAM configuration structure is presented in Fig. 2. The main MEF-OAM configuration guidelines are the following:

- When configuring an OAM bridge, the administrators have to always specify the VLAN ID on which information is gathered. Up to five VLAN IDs and seven MEF-OAM domains can be configured at the same time.
- When data collection is completed, the administrators can remove the OAM configuration.
- MEF-OAM supports a jitter accuracy of 5 milliseconds.

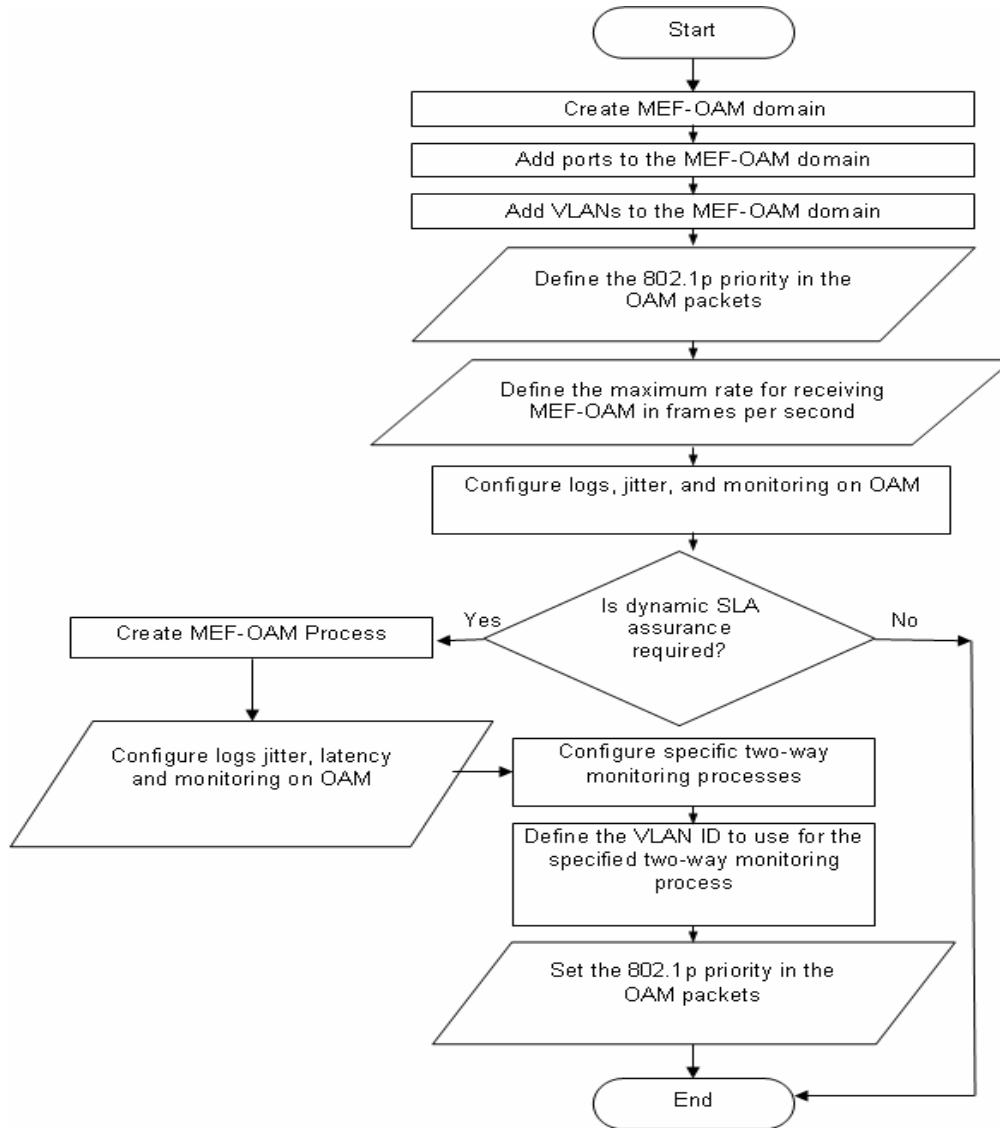


Fig. 2. MEF-OAM configuration

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], you gain broadcast Internet access 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) 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 [2].

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)

Standard EFM-OAM Structure

In Fig. 3 is presented the current EFM-OAM (also called 802.3ah Standard) module structure. It consists of three tasks that access the internal database together with the CLI task.

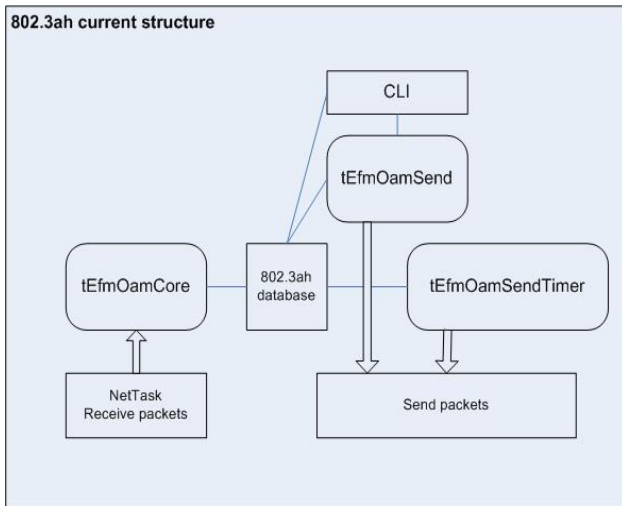


Fig. 3. Standard EFM-OAM structure

New Proposed EFM-OAM Structure

The EFM-OAM new structure is presented in Fig. 4. This structure has only one task **tEfmOamMain** (see Fig. 4) that accesses the internal database for a semaphore synchronization. The user (CLI or SNMP) 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.

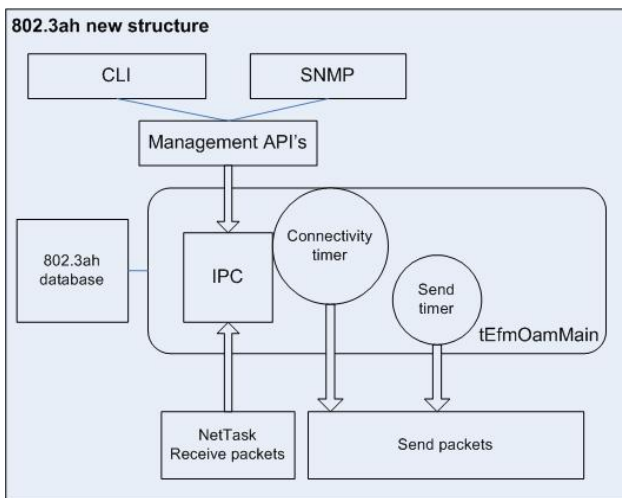


Fig. 4. New EFM-OAM structure

Connectivity timer:

- Replaces the **tEfmOamSendTimer** (see Fig. 3) 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).

Send timer:

- Replaces the **tEfmOamSend** (see Fig. 3) 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).

Potential Applications

Service providers use link layer EFM for demarcation point OAM services.

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.

EFM-OAM Protocol Functionality

EFM-OAM supports the following functionalities [10]:

- *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.
- *MIB variable retrieval*: retrieving information from a Management Information Base (MIB).
- *Organizing specific enhancements*: providing vendor-specific enhancements to the protocol.

Timers

Two configurable timers control the protocol:

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

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

Flags

Each OAMPDU includes a *Flags* field that includes the discovery process status. There are three possible status values [10]:

- *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.

5. CONCLUDING REMARKS

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 [7].

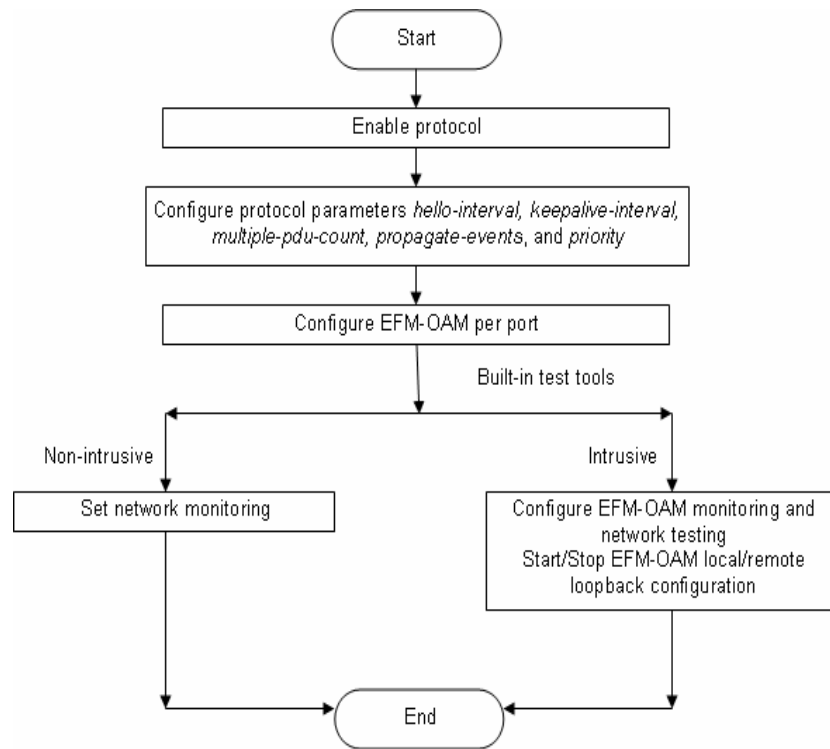


Fig. 5. EFM-OAM configuration

This implementation assures that OAMPDUs only affect the operation of the OAM protocol itself and not user data traffic.

Two protocols were designed to base on OAM feature: MEF and EFM. Their structure was modified as follows:

- Standard MEF protocol is based now on OAM.
- Standard EFM protocol has a new structure modified to work concomitant with MEF.

REFERENCES

- [1] Beck, M. (2005). *Ethernet in the First Mile: The IEEE802.3ah EFM Standard*, McGraw-Hill Professional.
- [2] Diab, W., Howard, M. F. (2006). *Ethernet in the First Mile: Access for Everyone*. IEEE Standards Information Network.
- [3] Iancu, E. (2004). *Data Transmission Theory*. Universitaria Publishing House, Craiova, Romania.
- [4] Iancu, E. and Vinătoru, M. (2003). *Analytical Methods for Detection and Localization of Defects. Case Studies*. Universitaria Publishing House, Craiova, Romania.
- [5] Karam, M. J. and Tobagi, F. A. (2002). *Analysis of delay and delay jitter of voice traffic in the internet*. The International Journal of Computer and Telecommunications Networking, vol. 40, no. 6.
- [6] ***http://www.cisco.com/en/US/prod/collateral/routers/ps368/prod_white_paper0900aecd804a0266.html
- [7] *** <http://www.faqs.org/rfcs/rfc4878.html>
- [8] *** <http://metroethernetforum.org/AbouttheMEF>, MEF FAQ Metro Ethernet Forum Membership Application.
- [9] http://metroethernetforum.org/MSWord_Documents/MEF10.doc.
- [10] *** <http://support.batm.com/products.php>
- [11] ***<http://whitepapers.zdnet.com/abstract.aspx?docid=952531&promo=100511&tag=content;coll>