

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 29, 2014

A. Takacs
Ericsson
D. Fedyk
Hewlett-Packard Company
J. He
Huawei
February 25, 2014

GMPLS RSVP-TE extensions for OAM Configuration
draft-ietf-ccamp-oam-configuration-fwk-13

Abstract

Operations, Administration and Maintenance (OAM) is an integral part of transport connections, hence it is required that OAM functions are activated/deactivated in sync with connection commissioning/decommissioning; avoiding spurious alarms and ensuring consistent operation. In certain technologies, OAM entities are inherently established once the connection is set up, while other technologies require extra configuration to establish and configure OAM entities. This document specifies extensions to RSVP-TE to support the establishment and configuration of OAM entities along with Label Switched Path signaling.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 29, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Requirements	4
3. RSVP-TE based OAM Configuration	6
3.1. Establishment of OAM Entities and Functions	8
3.2. Adjustment of OAM Parameters	9
3.3. Deleting OAM Entities	10
4. RSVP-TE Extensions	11
4.1. LSP Attributes Flags	11
4.2. OAM Configuration TLV	12
4.2.1. OAM Function Flags Sub-TLV	14
4.2.2. Technology Specific Sub-TLVs	14
4.3. Administrative Status Information	15
4.4. Handling OAM Configuration Errors	15
4.5. Considerations on Point-to-Multipoint OAM Configuration .	16
5. IANA Considerations	17
5.1. ADMIN_STATUS Object Bit Flags	17
5.2. LSP Attributes Flags	17
5.3. New LSP Attributes	18
5.4. RSVP Error Code	18
5.5. RSVP-TE OAM Configuration Registry	19
5.5.1. OAM Types Sub-Registry	19
5.5.2. OAM Sub-TLVs Sub-Registry	19
5.5.3. OAM Function Flags Sub-Registry	20
6. Security Considerations	20
7. Acknowledgements	21
8. References	21
8.1. Normative References	21
8.2. Informative References	21
Authors' Addresses	22

1. Introduction

GMPLS is designed as an out-of-band control plane supporting dynamic connection provisioning for any suitable data plane technology; including spatial switching (e.g., incoming port or fiber to outgoing port or fiber), wavelength-division multiplexing (e.g., Dense Wavelength Division Multiplexing (DWDM)), time-division multiplexing (e.g., SONET/SDH, G.709), and Ethernet Provider Backbone Bridging - Traffic Engineering (PBB-TE) and MPLS. In most of these technologies, there are Operations, Administration and Maintenance (OAM) functions employed to monitor the health and performance of the connections and to trigger data plane (DP) recovery mechanisms. Similar to connection provisioning, OAM functions follow general principles, but also have some technology specific characteristics.

OAM is an integral part of transport connections. Therefore it is required that OAM functions are activated/deactivated in sync with connection commissioning/decommissioning; avoiding spurious alarms and ensuring consistent operation. In certain technologies, OAM entities are inherently established once the connection is set up, while other technologies require extra configuration to establish and configure OAM entities. In some situations the use of OAM functions, such as Fault Management (FM) and Performance Management (PM), may be optional (based on network management policies). Hence, the network operator must be able to choose which set of OAM functions to apply to specific connections and which parameters should be configured and activated. To achieve this objective, OAM entities and specific functions must be selectively configurable.

In general, it is required that the management plane and control plane connection establishment mechanisms are synchronized with OAM establishment and activation. In particular, if the GMPLS control plane is employed, it is desirable to bind OAM setup and configuration to connection establishment signaling to avoid two separate management/configuration steps (connection setup followed by OAM configuration) which increases delay, processing, and more importantly may be prone to misconfiguration errors. Once OAM entities are setup and configured, pro-active as well as on-demand OAM functions can be activated via the management plane. On the other hand, it should be possible to activate/deactivate pro-active OAM functions via the GMPLS control plane as well. In some situations it may be possible to use the GMPLS control plane to control on-demand OAM functions too.

This document describes requirements for OAM configuration and control via Resource Reservation Protocol - Traffic Engineering (RSVP-TE). Extensions to the RSVP-TE protocol are specified providing a framework to configure and control OAM entities along

with the capability to carry technology specific information. Extensions can be grouped into: generic elements that are applicable to any OAM solution; and technology specific elements that provide additional configuration parameters, which may only be needed for a specific OAM technology. This document specifies the technology agnostic elements and specifies the way additional technology specific OAM parameters are provided.

This document addresses end-to-end OAM configuration, that is, the setup of OAM entities bound to an end-to-end Label-Switched Path (LSP), and configuration and control of OAM functions running end-to-end in the LSP. Configuration of OAM entities for LSP segments and tandem connections are out of the scope of this document.

The mechanisms described in this document provide an additional option for bootstrapping OAM that is not intended to replace or deprecate the use of other technology specific OAM bootstrapping techniques; e.g., LSP Ping [RFC4379] for MPLS networks. The procedures specified in this document are intended only for use in environments where RSVP-TE signaling is used to set up the LSPs that are to be monitored using OAM.

2. Requirements

This section summarizes various technology-specific OAM requirements which can be used as a basis for an OAM configuration framework.

MPLS OAM requirements are described in [RFC4377], which provides requirements to create consistent OAM functionality for MPLS networks. The following list is an excerpt of MPLS OAM requirements documented in [RFC4377] that bear a direct relevance to the discussion set forth in this document.

- o It is desired to support the automation of LSP defect detection. It is especially important in cases where large numbers of LSPs might be tested.
- o In particular some LSPs may require automated ingress-LSR to egress-LSR testing functionality, while others may not.
- o Mechanisms are required to coordinate network responses to defects. Such mechanisms may include alarm suppression, translating defect signals at technology boundaries, and synchronizing defect detection times by setting appropriately bounded detection time frames.

MPLS-TP defines a profile of MPLS targeted at transport applications [RFC5921]. This profile specifies the specific MPLS characteristics

and extensions required to meet transport requirements, including providing additional OAM, survivability, and other maintenance functions not currently supported by MPLS. Specific OAM requirements for MPLS-TP are specified in [RFC5654] and [RFC5860]. MPLS-TP poses the following requirements on the control plane to configure and control OAM entities:

- o From [RFC5860]: OAM functions MUST operate and be configurable even in the absence of a control plane. Conversely, it SHOULD be possible to configure as well as enable/disable the capability to operate OAM functions as part of connectivity management, and it SHOULD also be possible to configure as well as enable/disable the capability to operate OAM functions after connectivity has been established.
- o From [RFC5654]: The MPLS-TP control plane MUST support the configuration and modification of OAM maintenance points as well as the activation/ deactivation of OAM when the transport path or transport service is established or modified.

Ethernet Connectivity Fault Management (CFM) defines an adjunct connectivity monitoring OAM flow to check the liveness of Ethernet networks [IEEE.802.1Q-2011]. With PBB-TE [IEEE.802.1Q-2011] Ethernet networks support explicitly-routed Ethernet connections. CFM can be used to track the liveness of PBB-TE connections and detect data plane failures. In the IETF, the GMPLS controlled Ethernet Label Switching (GELS) (see [RFC5828] and [RFC6060]) work extended the GMPLS control plane to support the establishment of PBB-TE data plane connections. Without control plane support separate management commands would be needed to configure and start CFM.

GMPLS based OAM configuration and control needs to provide a general framework to be applicable to a wide range of data plane technologies and OAM solutions. There are three typical data plane technologies used for transport applications: wavelength based such as Wavelength Switched Optical Networks (WSON), Time-Division Multiplexing (TDM) based such as Synchronous Digital Hierarchy (SDH) and Synchronous Optical Networking (SONET), and packet based such as MPLS-TP [RFC5921] and Ethernet PBB-TE [IEEE.802.1Q-2011]. For all these data planes, the operator MUST be able to configure and control the following OAM functions:

- o It MUST be possible to explicitly request the setup of OAM entities for the signaled LSP and provide specific information for the setup if this is required by the technology.
- o Control of alarms is important to avoid false alarm indications and reporting to the management system. It MUST be possible to

enable/disable alarms generated by OAM functions. In some cases, selective alarm control may be desirable when, for instance, the operator is only concerned about critical alarms. Therefore the non-service affecting alarms should be inhibited.

- o When periodic messages are used for liveness check (continuity check) of LSPs, it MUST be possible to set the frequency of messages. This allows proper configuration for fulfilling the requirements of the service and/or meeting the detection time boundaries posed by possible congruent connectivity check operations of higher layer applications. For a network operator to be able to balance the trade-off between fast failure detection and data overhead, it is beneficial to configure the frequency of continuity check messages on a per LSP basis.
- o Pro-active Performance Monitoring (PM) functions are used to continuously collect information about specific characteristics of the connection. For consistent measurement of Service Level Agreements (SLAs), it MUST be possible to set common configuration parameters for the LSP.
- o The extensions MUST allow the operator to use only a minimal set of OAM configuration and control features if supported by the OAM solution or network management policy. Generic OAM parameters and data plane or OAM technology specific parameters MUST be supported.

3. RSVP-TE based OAM Configuration

In general, two types of maintenance points can be distinguished: Maintenance Entity Group End Points (MEPs) and Maintenance Entity Group Intermediate Points (MIPs). MEPs reside at the ends of an LSP and are capable of initiating and terminating OAM messages for Fault Management (FM) and Performance Monitoring (PM). MIPs on the other hand, are located at transit nodes of an LSP and are capable of reacting to some OAM messages but otherwise do not initiate messages. Maintenance Entity (ME) refers to an association of MEPs and MIPs that are provisioned to monitor an LSP.

When an LSP is signaled, a forwarding association is established between endpoints and transit nodes via label bindings. This association creates a context for the OAM entities monitoring the LSP. On top of this association, OAM entities may be configured to unambiguously identify MEs.

In addition to ME identification parameters, pro-active OAM functions (e.g., Continuity Check (CC) and Performance Monitoring (PM)) may have additional parameters that require configuration as well. In

particular, the frequency of periodic CC packets and the measurement interval for loss and delay measurements may need to be configured.

The above parameters may be either derived from LSP provisioning information, or alternatively, pre-configured default values can be used. In the simplest case, the control plane MAY provide information on whether or not OAM entities need to be setup for the signaled LSP. If OAM entities are created, control plane signaling MUST also provide a means to activate/deactivate OAM message flows and associated alarms.

OAM identifiers, as well as the configuration of OAM functions, are technology specific (i.e., vary depending on the data plane technology and the chosen OAM solution). In addition, for any given data plane technology, a set of OAM solutions may be applicable. Therefore, the OAM configuration framework allows selecting a specific OAM solution to be used for the signaled LSP and provides means to carry detailed OAM configuration information in technology specific TLVs.

Administrative Status Information is carried in the ADMIN_STATUS Object. The Administrative Status Information is described in [RFC3471], the ADMIN_STATUS Object is specified for RSVP-TE in [RFC3473]. Two bits are allocated for the administrative control of OAM monitoring: the "OAM Flows Enabled" (M) and "OAM Alarms Enabled" (O) bits. When the "OAM Flows Enabled" bit is set, OAM mechanisms MUST be enabled; if it is cleared, OAM mechanisms MUST be disabled. When the "OAM Alarms Enabled" bit is set OAM triggered alarms are enabled and associated consequent actions MUST be executed including the notification to the management system. When this bit is cleared, alarms are suppressed and no action SHOULD be executed and the management system SHOULD NOT be notified.

The LSP_ATTRIBUTES and the LSP_REQUIRED_ATTRIBUTES objects are defined in [RFC5420] to provide means to signal LSP attributes and options in the form of TLVs. Options and attributes signaled in the LSP_ATTRIBUTES object can be passed transparently through LSRs not supporting a particular option or attribute, while the contents of the LSP_REQUIRED_ATTRIBUTES object MUST be examined and processed by each LSR. One bit "OAM MEP entities desired" is allocated in the LSP Attributes Flags TLV to be used in the LSP_ATTRIBUTES object. If the "OAM MEP entities desired" bit is set it is indicating that the establishment of OAM MEP entities are required at the endpoints of the signaled LSP. One bit "OAM MIP entities desired" is allocated in the LSP Attributes Flags TLV to be used in the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES objects. If the "OAM MIP entities desired" bit is set in the LSP_ATTRIBUTES Flags TLV in the LSP_REQUIRED_ATTRIBUTES Object, it is indicating that the

establishment of OAM MIP entities is required at every transit node of the signaled LSP.

3.1. Establishment of OAM Entities and Functions

In order to avoid spurious alarms, OAM functions should be setup and enabled in the appropriate order. When using the GMPLS control plane for both LSP establishment and to enable OAM functions on the LSPs, the control of both processes is bound to RSVP-TE message exchanges.

An LSP may be signaled and established without OAM configuration first, and OAM entities may be added later with a subsequent re-signaling of the LSP. Alternatively, the LSP may be setup with OAM entities with the first signaling of the LSP. The below procedures apply to both cases.

Before initiating a Path message with OAM Configuration information, an initiating node MUST establish and configure the corresponding OAM entities locally. But until the LSP is established, OAM source functions MUST NOT start sending any OAM messages. In the case of bidirectional connections, in addition to the OAM source function, the initiator node MUST set up the OAM sink function and prepare it to receive OAM messages. During this time the OAM alarms MUST be suppressed (e.g., due to missing or unidentified OAM messages). To achieve OAM alarm suppression, Path message MUST be sent with the "OAM Alarms Enabled" ADMIN_STATUS flag cleared.

When the Path message arrives at the receiver, the remote end MUST establish and configure OAM entities according to the OAM information provided in the Path message. If this is not possible, a PathErr message SHOULD be sent and neither the OAM entities nor the LSP SHOULD be established. If OAM entities are established successfully, the OAM sink function MUST be prepared to receive OAM messages, but MUST NOT generate any OAM alarms (e.g., due to missing or unidentified OAM messages). In the case of bidirectional connections, in addition to the OAM sink function, an OAM source function MUST be set up and, according to the requested configuration, the OAM source function MUST start sending OAM messages. Then a Resv message MUST be sent back, including the LSP_Attributes Flags TLV, with the appropriate setting of the "OAM MEP entities desired" and "OAM MIP entities desired" flags, and the OAM Configuration TLV that corresponds to the established and configured OAM entities and functions. Depending on the OAM technology, some elements of the OAM Configuration TLV MAY be updated /changed; i.e., if the remote end is not supporting a certain OAM configuration it may suggest an alternative setting, which may or may not be accepted by the initiator of the Path message. If it is accepted, the initiator will reconfigure its OAM functions according

to the information received in the Resv message. If the alternate setting is not acceptable a ResvErr message MAY be sent tearing down the LSP. Details of this operation are technology specific and should be described in accompanying technology specific documents.

When the initiating side receives the Resv message, it completes any pending OAM configuration and enables the OAM source function to send OAM messages.

After this exchange, OAM entities are established and configured for the LSP and OAM messages are exchanged. OAM alarms can now be enabled. The initiator, during the period when OAM alarms are disabled, sends a Path message with "OAM Alarms Enabled" ADMIN_STATUS flag set. The receiving node enables the OAM alarms after processing the Path message. The initiator enables OAM alarms after it receives the Resv message. Data plane OAM is now fully functional.

In case an egress LSR does not support the extensions defined in this document, according to [RFC5420], it will silently ignore the new LSP Attributes Flags as well as the TLVs carrying additional OAM configuration information, and therefore no error will be raised that would notify the ingress LSR about the missing OAM configuration actions on the egress side. However, as described above, an egress LSR conformant to the specification of this document will set the LSP Attributes Flags and include the OAM Configuration TLV in the Resv message indicating the configuration of the OAM mechanisms, therefore an ingress LSR by detecting the missing information in the Resv message will be able to recognize that the remote end does not support the OAM configuration functionality and therefore it SHOULD tear down the LSP, and if appropriate, signal the LSP without any OAM configuration information.

3.2. Adjustment of OAM Parameters

There may be a need to change the parameters of an already established and configured OAM function during the lifetime of the LSP. To do so the LSP needs to be re-sigaled with the updated parameters. OAM parameters influence the content and timing of OAM messages and identify the way OAM defects and alarms are derived and generated. Hence, to avoid spurious alarms, it is important that both sides, OAM sink and source, are updated in a synchronized way. First, the alarms of the OAM sink function should be suppressed and only then should expected OAM parameters be adjusted. Subsequently, the parameters of the OAM source function can be updated. Finally, the alarms of the OAM sink side can be enabled again.

In accordance with the above operation, the LSP MUST first be re-sigaled with "OAM Alarms Enabled" ADMIN_STATUS flag cleared,

including the updated OAM Configuration TLV corresponding to the new parameter settings. The initiator MUST keep its OAM sink and source functions running unmodified, but it MUST suppress OAM alarms after the updated Path message is sent. The receiver MUST first disable all OAM alarms, then update the OAM parameters according to the information in the Path message and reply with a Resv message acknowledging the changes by including the OAM Configuration TLV. Note that the receiving side has the possibility to adjust the requested OAM configuration parameters and reply with an updated OAM Configuration TLV in the Resv message, reflecting the actually configured values. However, in order to avoid an extensive negotiation phase, in the case of adjusting already configured OAM functions, the receiving side SHOULD NOT update the parameters requested in the Path message to an extent that would provide lower performance (e.g., lower frequency of monitoring packets) than what has been in operation previously.

The initiator MUST only update its OAM sink and source functions after it received the Resv message. After this Path/Resv message exchange (in both unidirectional and bidirectional LSP cases) the OAM parameters are updated and OAM is running according the new parameter settings. However, OAM alarms are still disabled. A subsequent Path /Resv message exchange with "OAM Alarms Enabled" ADMIN_STATUS flag set is needed to enable OAM alarms again.

3.3. Deleting OAM Entities

In some cases it may be useful to remove some or all OAM entities and functions from an LSP without actually tearing down the connection.

To avoid any spurious alarms, first the LSP MUST be re-signaled with "OAM Alarms Enabled" ADMIN_STATUS flag cleared but unchanged OAM configuration. Subsequently, the LSP is re-signaled with "OAM MEP Entities desired" and "OAM MIP Entities desired" LSP ATTRIBUTES flags cleared, and without the OAM Configuration TLV, this MUST result in the deletion of all OAM entities associated with the LSP. All control and data plane resources in use by the OAM entities and functions SHOULD be freed up. Alternatively, if only some OAM functions need to be removed, the LSP is re-signaled with the updated OAM Configuration TLV. Changes between the contents of the previously signaled OAM Configuration TLV and the currently received TLV represent which functions MUST be removed/added.

OAM source functions MUST be deleted first and only after the "OAM Alarms Disabled" can the associated OAM sink functions be removed, this will ensure that OAM messages do not leak outside the LSP. To this end the initiator, before sending the Path message, MUST remove the OAM source, hence terminating the OAM message flow associated to

the downstream direction. In the case of a bidirectional connection, it MUST leave in place the OAM sink functions associated to the upstream direction. The remote end, after receiving the Path message, MUST remove all associated OAM entities and functions and reply with a Resv message without an OAM Configuration TLV. The initiator completely removes OAM entities and functions after the Resv message arrived.

4. RSVP-TE Extensions

RFC Editor Note: remove/update "IANA" and "IANA to assign" notes in the document once the assignments have been made.

4.1. LSP Attributes Flags

In RSVP-TE the Flags field of the SESSION_ATTRIBUTE object is used to indicate options and attributes of the LSP. The Flags field has 8 bits and hence is limited to differentiate only 8 options. [RFC5420] defines new objects for RSVP-TE messages to allow the signaling of arbitrary attribute parameters making RSVP-TE easily extensible to support new applications. Furthermore, [RFC5420] allows options and attributes that do not need to be acted on by all Label Switched Routers (LSRs) along the path of the LSP. In particular, these options and attributes may apply only to key LSRs on the path such as the ingress LSR and egress LSR. Options and attributes can be signaled transparently, and only examined at those points that need to act on them. The LSP_ATTRIBUTES and the LSP_REQUIRED_ATTRIBUTES objects are defined in [RFC5420] to provide means to signal LSP attributes and options in the form of TLVs. Options and attributes signaled in the LSP_ATTRIBUTES object can be passed transparently through LSRs not supporting a particular option or attribute, while the contents of the LSP_REQUIRED_ATTRIBUTES object MUST be examined and processed by each LSR. One TLV is defined in [RFC5420]: the Attributes Flags TLV.


One bit (IANA to assign): "OAM MEP entities desired" is allocated in the LSP Attributes Flags TLV to be used in the LSP_ATTRIBUTES object. If the "OAM MEP entities desired" bit is set it is indicating that the establishment of OAM MEP entities are required at the endpoints of the signaled LSP. If the establishment of MEPs is not supported an error MUST be generated: "OAM Problem/MEP establishment not supported".

If the "OAM MEP entities desired" bit is set and additional parameters need to be configured, an OAM Configuration TLV MAY be included in the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES object.

One bit (IANA to assign): "OAM MIP entities desired" is allocated in the LSP Attributes Flags TLV to be used in the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES objects. If the "OAM MEP entities desired" bit is not set then this bit MUST NOT be set. If the "OAM MIP entities desired" bit is set in the LSP_ATTRIBUTES Flags TLV in the LSP_REQUIRED_ATTRIBUTES Object, it is indicating that the establishment of OAM MIP entities is required at every transit node of the signaled LSP. If the establishment of a MIP is not supported an error MUST be generated: "OAM Problem/MIP establishment not supported". If an intermediate LSR does not support the extensions defined in this document it will not recognize the "OAM MIP entities desired" flag and although the LSP_REQUIRED_ATTRIBUTES object was used it will not configure MIP entities and will not raise any errors. If LSRs that are not supporting this document are to be assumed in the network, the ingress LSR SHOULD collect per-hop information about the LSP Attributes utilizing the LSP Attributes sub-object of the Record Route Object as defined in [RFC5420]. When the Record Route object is received the ingress SHOULD check whether all intermediate LSRs set the "OAM MIP entities desired" flag indicating support of the function, if not, depending on operator policy the LSP MAY need to be torn down.

4.2. OAM Configuration TLV

This TLV provides information about which OAM technology/method should be used and carries sub-TLVs for any additional OAM configuration information. One OAM Configuration TLV MAY be carried in the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES object in Path and Resv messages. When carried in the LSP_REQUIRED_ATTRIBUTES object, it is indicating that intermediate nodes MUST recognize and react on the OAM configuration information.



Type: indicates a new type: the OAM Configuration TLV (IANA to assign).

OAM Type: specifies the technology specific OAM method. When carried in the LSP_REQUIRED_ATTRIBUTES Object, if the requested OAM method is not supported at any given node an error MUST be generated: "OAM Problem/Unsupported OAM Type". When carried in the LSP_ATTRIBUTES Object, intermediate nodes not supporting the OAM Type pass the object forward unchanged as specified in [RFC5420]. Ingress and egress nodes that support the OAM Configuration TLV but that do not support a specific OAM Type MUST respond with an error indicating "OAM Problem/Unsupported OAM Type".

OAM Type	Description
-----	-----
0-255	Reserved

This document defines no types. IANA is requested to maintain the values in a new "RSVP-TE OAM Configuration Registry".

Length: indicates the total length of the TLV in octets. The TLV MUST be zero-padded so that the TLV is four octet aligned.

Two groups of TLVs are defined: generic sub-TLVs and technology specific sub-TLVs. Generic sub-TLVs carry information that are applicable independent of the actual OAM technology, while technology specific sub-TLVs are providing configuration parameters for specific OAM technologies. This document defines one generic sub-TLV, see Section 4.2.1, while it is foreseen that technology specific sub-TLVs will be defined by separate documents.

The receiving node, based on the OAM Type, will check if a corresponding technology specific OAM configuration sub-TLV is included in the OAM Configuration TLV. If the included technology specific OAM configuration sub-TLV is different from what is specified in the OAM Type an error MUST be generated: "OAM Problem/OAM Type Mismatch". IANA is requested to maintain the sub-TLV space in the new "RSVP-TE OAM Configuration Registry".


Sub-TLV Type	Description
-----	-----
0	Reserved
1	OAM Function Flags Sub-TLV
2-31	Reserved for generic Sub-TLVs
32-	Reserved for technology specific Sub-TLVs

Note that there is a hierarchical dependency between the OAM configuration elements. First, the "OAM MEP entities desired" flag needs to be set. Only when that flag is set MAY an "OAM Configuration TLV" be included in the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES Object. When this TLV is present, based on

the "OAM Type" field, it MAY carry a technology specific OAM configuration sub-TLV. If this hierarchy is broken (e.g., "OAM MEP entities desired" flag is not set but an OAM Configuration TLV is present) an error MUST be generated: "OAM Problem/Configuration Error".

4.2.1. OAM Function Flags Sub-TLV

The "OAM Configuration TLV" MUST always include a single instance of the "OAM Function Flags Sub-TLV" and it MUST always be the first sub-TLV. "OAM Function Flags" specifies which pro-active OAM functions (e.g., connectivity monitoring, loss and delay measurement) and which fault management signals MUST be established and configured. If the selected OAM Function(s) is(are) not supported, an error MUST be generated: "OAM Problem/Unsupported OAM Function".



OAM Function Flags is bitmap with extensible length based on the Length field of the TLV. Bits are numbered from left to right. The TLV is padded to 4-octet alignment. The Length field indicates the size of the padded TLV in octets. IANA is requested to maintain the OAM Function Flags in the new "RSVP-TE OAM Configuration Registry". This document defines the following flags.

OAM Function Flag bit#	Description
0	Continuity Check (CC)
1	Connectivity Verification (CV)
2	Fault Management Signal (FMS)
3	Performance Monitoring/Loss (PM/Loss)
4	Performance Monitoring/Delay (PM/Delay)
5	Performance Monitoring/Throughput Measurement (PM/Throughput)

4.2.2. Technology Specific Sub-TLVs

If technology-specific configuration information is needed for a specific "OAM Type", then this information is carried in a technology-specific sub-TLV. Such sub-TLVs are OPTIONAL and an OAM

Configuration TLV MUST NOT contain more than one technology- specific sub-TLV. IANA is requested to maintain the OAM technology specific sub-TLV space in the new "RSVP-TE OAM Configuration Registry".

4.3. Administrative Status Information

Administrative Status Information is carried in the ADMIN_STATUS Object, which is specified for RSVP-TE in [RFC3473]. The Administrative Status Information is described in [RFC3471].

Two bits are allocated for the administrative control of OAM monitoring. Two bits (IANA to assign) are allocated by this document: the "OAM Flows Enabled" (M) and "OAM Alarms Enabled" (O) bits. When the "OAM Flows Enabled" bit is set, OAM mechanisms MUST be enabled; if it is cleared, OAM mechanisms MUST be disabled. When the "OAM Alarms Enabled" bit is set OAM triggered alarms are enabled and associated consequent actions MUST be executed including the notification to the management system. When this bit is cleared, alarms are suppressed and no action SHOULD be executed and the management system SHOULD NOT be notified. For a detailed description of the use of these flags see Section 3.

4.4. Handling OAM Configuration Errors

To handle OAM configuration errors, a new Error Code (IANA to assign) "OAM Problem" is introduced. To refer to specific problems, a set of Error Values are defined under the "OAM Problem" error code.

If a node does not support the establishment of OAM MEP or MIP entities it MUST use the error value: "MEP establishment not supported" or "MIP establishment not supported" respectively in the PathErr message.

If a node does not support a specific OAM technology/solution it MUST use the error value: "Unsupported OAM Type" in the PathErr message.

If a different technology specific OAM configuration TLV is included than what was specified in the OAM Type an error MUST be generated with error value: "OAM Type Mismatch" in the PathErr message.

There is a hierarchy between the OAM configuration elements. If this hierarchy is broken, the error value: "Configuration Error" MUST be used in the PathErr message.

If a node does not support a specific OAM Function, it MUST use the error value: "Unsupported OAM Function" in the PathErr message.

4.5. Considerations on Point-to-Multipoint OAM Configuration

RSVP-TE extensions for the establishment of point-to-multipoint (P2MP) LSPs are specified in [RFC4875]. A P2MP LSP is comprised of multiple source-to-leaf (S2L) sub-LSPs. These S2L sub-LSPs are set up between the ingress and egress LSRs, and are appropriately combined by the branch LSRs using RSVP semantics to result in a P2MP TE LSP. One Path message may signal one or multiple S2L sub-LSPs for a single P2MP LSP. Hence, the S2L sub-LSPs belonging to a P2MP LSP can be signaled using one Path message or split across multiple Path messages.

P2MP OAM mechanisms are very specific to the data plane technology, therefore in this document, we only highlight the basic principles of P2MP OAM configuration. We consider only the root to leaf OAM flows, and as such, aspects of the configuration of return paths are outside the scope of our discussions. We also limit our consideration to the case where all leaves must successfully establish OAM entities with identical configuration in order the P2MP OAM is successfully established. In any case, the discussion set forth below provides only guidelines for P2MP OAM configuration. However at minimum the below procedures SHOULD be specified for P2MP OAM configuration in a technology specific document.

The root node may use a single Path message or multiple Path messages to setup the whole P2MP tree. In the case when multiple Path messages are used, the root node is responsible to keep the OAM Configuration information consistent in each of the sent Path messages, i.e., the same information MUST be included in all Path messages used to construct the multicast tree. Each branching node will propagate the Path message downstream on each of the branches, when constructing a Path message the OAM Configuration information MUST be copied unchanged from the received Path message, including the related ADMIN_STATUS bits, LSP Attribute Flags and the OAM Configuration TLV. The latter two also imply that the LSP_ATTRIBUTES and LSP_REQUIRED_ATTRIBUTES Object MUST be copied for the upstream Path message to the subsequent downstream Path messages.

Leaves MUST create and configure OAM sink functions according to the parameters received in the Path message, for P2MP OAM configuration there is no possibility for parameter negotiation on a per leaf basis. This is due to the fact that the OAM source function, residing in the root of the tree, will operate with a single configuration, which then must be obeyed by all leaves. If a leaf cannot accept the OAM parameters it MUST use the RRO Attributes sub-object [RFC5420] to notify the root about the problem. In particular, if the OAM configuration was successful, the leaf would set the "OAM MEP entities desired" flag in the RRO Attributes sub-

object in the Resv message. On the other hand, if OAM entities could not be established the Resv message should be sent with the "OAM MEP entities desired" bit cleared in the RRO Attributes sub-object. Branching nodes should collect and merge the received RROs according to the procedures described in [RFC4875]. This way, the root when receiving the Resv message (or messages if multiple Path messages were used to set up the tree) will have a clear information about which of the leaves could establish the OAM functions. If all leaves established OAM entities successfully, the root can enable the OAM message flow. On the other hand, if at some leaves the establishment was unsuccessful additional actions will be needed before the OAM message flow can be enabled. Such action could be to setup two independent P2MP LSPs. One LSP with OAM Configuration information towards leaves which could successfully setup the OAM function. This can be done by pruning the leaves which failed to setup OAM of the previously signaled P2MP LSP. The other P2MP LSP could be constructed for leaves without OAM entities. The exact procedures will be described in technology specific documents.

5. IANA Considerations

5.1. ADMIN_STATUS Object Bit Flags

IANA maintains a registry called "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters" with a sub-registry called "Administrative Status Information Flags".

IANA is requested to allocate two new flags as follows:

Bit Number	Hex Value	Name	Reference
TBA	TBA	OAM Alarms Enabled (O)	[This.ID]
TBA	TBA	OAM Flows Enabled (M)	[This.ID]

5.2. LSP Attributes Flags

IANA maintains a registry called "Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Parameters" with a subregistry called "Attribute Flags".

IANA is requested to allocate two new flags as follows:

Bit No	Name	Attribute Flags Path	Attribute Flags Resv	RRO	Reference
TBA	OAM MEP entities desired	Yes	Yes	Yes	[This.ID]
TBA	OAM MIP entities desired	Yes	Yes	Yes	[This.ID]

5.3. New LSP Attributes

IANA maintains a registry called "Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Parameters" with a subregistry called "Attributes TLV Space"

IANA is requested to allocate one new TLV type as follows:

Type	Name	Allowed on LSP_ATTRIBUTES	Allowed on LSP_REQUIRED_ATTRIBUTES	Reference
TBA	OAM Configuration TLV	Yes	Yes	[This.ID]

5.4. RSVP Error Code

IANA maintains a registry called "Resource Reservation Protocol (RSVP) Parameters" with a subregistry called "Error Codes and Globally-Defined Error Value Sub-Codes".

IANA is requested to allocate one new Error Code as follows:

Error Code	Meaning	Reference
TBA	OAM Problem	[This.ID]

The value is to be selected from the range 0-239.

The following Error Value sub-codes are defined for this new Error Code as follows:

Value	Description	Reference
0	Reserved	[This.ID]
1	MEP establishment not supported	[This.ID]
2	MIP establishment not supported	[This.ID]
3	Unsupported OAM Type	[This.ID]
4	Configuration Error	[This.ID]
5	OAM Type Mismatch	[This.ID]
6	Unsupported OAM Function	[This.ID]
7-32767	Unassigned	
32768-65535	Reserved for Private Use	[This.ID]

5.5. RSVP-TE OAM Configuration Registry

IANA is requested to create a new registry called "RSVP-TE OAM Configuration Registry".

IANA is requested to create sub-registries as defined in the following subsections. The registration procedures specified are as defined in [RFC5226].

5.5.1. OAM Types Sub-Registry

IANA is requested to create the "OAM Types" sub-registry of the "RSVP-TE OAM Configuration Registry" as follows:

Range	Registration Procedures
0-255	IETF Review

There are no initial values in this registry. IANA should show the registry as follows:

OAM Type Number	OAM Type Description	Reference
0-255	Unassigned	

5.5.2. OAM Sub-TLVs Sub-Registry

IANA is requested to create the "OAM Sub-TLVs" sub-registry of the "RSVP-TE OAM Configuration Registry" as follows:

Range	Purpose	Registration Procedures
0-31	Generic Sub-TLVs	IETF Review
32-65534	Technology-specific Sub-TLVs	IETF Review
65535-65536	Experimental Sub-TLVs	Experimental

IANA is requested to populate the registry as follows:

Sub-TLV Type	Description	Reference
0	Reserved	[This.ID]
1	OAM Function Flags Sub-TLV	[This.ID]
2-31	Unassigned	
32-65534	Unassigned	

5.5.3. OAM Function Flags Sub-Registry

IANA is requested to create the "OAM Function Flags Sub-Registry" sub-registry of the "RSVP-TE OAM Configuration Registry".

New values in the registry are allocated by "IETF Review". There is no top value to the range. Bits are counted from bit 0 as the first bit transmitted.

IANA is requested to populate the registry as follows.

OAM Function Flag bit number	Description
0	Continuity Check (CC)
1	Connectivity Verification (CV)
2	Fault Management Signal (FMS)
3	Performance Monitoring/Loss (PM/Loss)
4	Performance Monitoring/Delay (PM/Delay)
5	Performance Monitoring/Throughput Measurement (PM/Throughput)
6-...	Unassigned

6. Security Considerations

The signaling of OAM related parameters and the automatic establishment of OAM entities based on RSVP-TE messages adds a new aspect to the security considerations discussed in [RFC3473]. In particular, a network element could be overloaded, if a remote attacker could request liveness monitoring, with frequent periodic messages, for a high number of LSPs, targeting a single network element. Such an attack can efficiently be prevented when mechanisms for message integrity and node authentication are deployed. Since

the OAM configuration extensions rely on the hop-by-hop exchange of exiting RSVP-TE messages, procedures specified for RSVP message security in [RFC2747] can be used to mitigate possible attacks.

For a more comprehensive discussion of GMPLS security, and attack mitigation techniques, please see the Security Framework for MPLS and GMPLS Networks [RFC5920].

7. Acknowledgements

The authors would like to thank Francesco Fondelli, Adrian Farrel, Loa Andersson, Eric Gray and Dimitri Papadimitriou for their useful comments.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5420] Farrel, A., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.

8.2. Informative References

- [IEEE.802.1Q-2011] IEEE, "IEEE Standard for Local and metropolitan area networks -- Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q, 2011.
- [RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.

- [RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and S. Matsushima, "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [RFC4875] Aggarwal, R., Papadimitriou, D., and S. Yasukawa, "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [RFC5654] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.
- [RFC5828] Fedyk, D., Berger, L., and L. Andersson, "Generalized Multiprotocol Label Switching (GMPLS) Ethernet Label Switching Architecture and Framework", RFC 5828, March 2010.
- [RFC5860] Vigoureux, M., Ward, D., and M. Betts, "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks", RFC 5860, May 2010.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [RFC5921] Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010.
- [RFC6060] Fedyk, D., Shah, H., Bitar, N., and A. Takacs, "Generalized Multiprotocol Label Switching (GMPLS) Control of Ethernet Provider Backbone Traffic Engineering (PBB-TE)", RFC 6060, March 2011.

Authors' Addresses

Attila Takacs
Ericsson
Konyves Kalman krt. 11.
Budapest 1097
Hungary

Email: attila.takacs@ericsson.com

Don Fedyk
Hewlett-Packard Company
153 Taylor Street
Littleton, MA 01460
USA

Email: don.fedyk@hp.com

Jia He
Huawei

Email: hejia@huawei.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 23, 2015

A. Takacs
B. Gero
Ericsson
H. Long
Huawei
July 22, 2014

GMPLS RSVP-TE Extensions for Ethernet OAM Configuration
draft-ietf-ccamp-rsvp-te-eth-oam-ext-13

Abstract

The GMPLS controlled Ethernet Label Switching (GELS) work extended GMPLS RSVP-TE to support the establishment of Ethernet LSPs. IEEE Ethernet Connectivity Fault Management (CFM) specifies an adjunct OAM flow to check connectivity in Ethernet networks. CFM can be also used with Ethernet LSPs for fault detection and triggering recovery mechanisms. The ITU-T Y.1731 specification builds on CFM and specifies additional OAM mechanisms, including Performance Monitoring, for Ethernet networks. This document specifies extensions of GMPLS RSVP-TE protocol to support the setup of the associated Ethernet OAM entities of Ethernet LSPs, and defines the Ethernet technology specific TLVs based on the GMPLS OAM Configuration Framework. This document supports, but does not modify, the IEEE and ITU-T OAM mechanisms.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 23, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Background	2
2. Overview of Ethernet OAM operation	3
3. GMPLS RSVP-TE Extensions	5
3.1. Operation overview	5
3.2. OAM Configuration TLV	7
3.3. Ethernet OAM Configuration Sub-TLV	8
3.3.1. MD Name Sub-TLV	8
3.3.2. Short MA Name Sub-TLV	9
3.3.3. MEP ID Sub-TLV	10
3.3.4. Continuity Check (CC) Sub-TLV	11
3.4. Pro-active Performance Monitoring	12
3.5. Summary of Ethernet OAM configuration errors	13
4. IANA Considerations	13
4.1. RSVP-TE OAM Configuration Registry	13
4.2. Ethernet Sub-TLVs Sub-Registry	14
4.3. RSVP Error Code	14
5. Security Considerations	15
6. Acknowledgements	15
7. Contributors	15
8. References	16
8.1. Normative References	16
8.2. Informative References	16
Authors' Addresses	16

1. Background

Provider Backbone Bridging - Traffic Engineering (PBB-TE) [IEEE.802.1Q-2011] decouples the Ethernet data and control planes, and allows external control and management mechanisms to create

explicitly routed Ethernet connections. In addition, PBB-TE defines mechanisms for protection switching of bidirectional Ethernet connections. Ethernet Connectivity Fault Management (CFM) defines an adjunct connectivity monitoring OAM flow to check the liveness of Ethernet networks [IEEE.802.1Q-2011], including the monitoring of specific explicitly routed Ethernet connections. The ITU-T Recommendation Y.1731 [ITU-T.Y.1731-2011] extended CFM and specified additional OAM functionality.

In IETF, the GMPLS controlled Ethernet Label Switching (GELS) work extended the GMPLS control plane to support the establishment of explicitly routed Ethernet connections [RFC5828][RFC6060]. We refer to GMPLS established Ethernet connections as Ethernet LSPs. GELS enables the application of MPLS-TE and GMPLS provisioning and recovery features in Ethernet networks.

The use of GMPLS RSVP-TE to support the establishment and configuration of OAM entities with LSP signaling is defined in a technology agnostic way in [RFC7260]. The purpose of this document is to specify the additional technology specific OAM entities to support Ethernet connections.

2. Overview of Ethernet OAM operation

For the purposes of this document, we only discuss Ethernet OAM aspects that are relevant for proactive connectivity monitoring of Ethernet LSPs. On-demand OAM functions for the purposes of this document will be supported by Management Plane operations.

PBB-TE defines point-to-point Ethernet Switched Paths (ESPs) as a provisioned traffic engineered unidirectional connectivity, identified by the 3-tuple [ESP-MAC DA, ESP-MAC SA, ESP-VID], where the ESP-MAC DA is the destination address of the ESP, the ESP-MAC SA is the source address of the ESP, and the ESP-VID is a VLAN identifier allocated for explicitly routed connections. To form a bidirectional PBB-TE connection, two co-routed point-to-point ESPs are combined. The combined ESPs must have the same ESP-MAC addresses but may have different ESP-VIDs. The formed co-routed bidirectional path is a path where the forward and backward directions follow the same route (links and nodes) across the network.

Note that although it would be possible to use GMPLS to setup a single unidirectional ESP, the Ethernet OAM mechanisms are only fully functional when bidirectional connections are established with co-routed ESPs. Therefore, the scope of this document only covers bidirectional point-to-point PBB-TE connections.

At both ends of the bidirectional point-to-point PBB-TE connection, one Maintenance Endpoint (MEP) is configured. The MEPs monitoring a PBB-TE connection must be configured with the same Maintenance Domain Level (MD Level) and Maintenance Association Identifier (MAID). Each MEP has a unique identifier, the MEP ID. Besides these identifiers, a MEP monitoring a PBB-TE connection must be provisioned with the 3-tuples [ESP-MAC DA, ESP-MAC SA, ESP-VID] of the two ESPs.

In the case of point-to-point VLAN connections, the connection may be identified with a single VLAN, or with two VLANs, one for each direction. Therefore, instead of the 3-tuples of the PBB-TE ESPs, MEPs must be provisioned with the proper VLAN identifiers.

MEPs exchange Connectivity Check Messages (CCMs) periodically with fixed intervals. Eight distinct intervals are defined in [IEEE.802.1Q-2011]:

#	CCM Interval (CCI)	3 bit encoding
0	Reserved	000
1	3 1/3 ms	001
2	10 ms	010
3	100 ms	011
4	1 s	100
5	10 s	101
6	1 min	110
7	10 min	111

Table 1: CCM Interval encoding

If 3 consecutive CCM messages are lost; connectivity failure is declared. The MEP detecting the failure will signal the defect to the remote MEP in the subsequent CCM messages it emits, by setting the Remote Defect Indicator (RDI) bit in the CCM message. If a MEP receives a CCM message with RDI bit set it immediately declares failure. The detection of a failure may trigger protection switching mechanisms or may be signaled to a management system.

At each transit node, Maintenance Intermediate Points (MIPs) may be established to help failure localization, e.g., using link trace and loop back functions. MIPs need to be provisioned with a subset of the MEP identification parameters described above.

3. GMPLS RSVP-TE Extensions

3.1. Operation overview

To simplify the configuration of connectivity monitoring, when an Ethernet LSP is signaled, the associated MEPs should be automatically established. To monitor an Ethernet LSP, a set of parameters must be provided to setup a Maintenance Association and related MEPs. Optionally, MIPs may be created at the transit nodes of the Ethernet LSP. The LSP Attribute Flags: "OAM MEP entities desired" and "OAM MIP entities desired", as described in [RFC7260], are used to signal that the respective OAM entities must be established. An OAM Configuration TLV, as described in [RFC7260], is added to the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES Objects specifying that Ethernet OAM is to be setup for the LSP. Ethernet OAM specific information, as described below, is carried in the new Ethernet OAM Configuration Sub-TLV (see Section 3.3) within the OAM Configuration TLV.

- o A unique MAID must be allocated for the PBB-TE connection and both MEPs must be configured with the same information. The MAID consists of an optional Maintenance Domain Name (MD Name) and a mandatory Short Maintenance Association Name (Short MA Name). Various formatting rules for these names have been defined in [IEEE.802.1Q-2011]. Since this information is also carried in all CCM messages, the combined length of the Names is limited to 44 bytes, see [IEEE.802.1Q-2011] for the details of the message format. How these parameters are determined is out of scope of this document.
- o Each MEP must be provisioned with a MEP ID. The MEP ID uniquely identifies a given MEP within a Maintenance Association. That is, the combination of MAID and MEP ID must uniquely identify a MEP. How the value of the MEP ID is determined is out of scope of this document.
- o The Maintenance Domain Level (MD Level) allows hierarchical separation of monitoring entities. [IEEE.802.1Q-2011] allows differentiation of 8 levels. How the value of the MD Level is determined is out of scope of this document. Note that probably for all Ethernet LSPs a single (default) MD Level will be used within a network domain.

- o The desired CCM Interval must be specified by the management system based on service requirements or operator policy. The same CCM Interval must be set in each of the MEPs monitoring a given Ethernet LSP. How the value of the CCM Interval is determined is out of scope of this document.
- o The desired forwarding priority to be set by MEPs for the CCM frames may be specified. The same CCM priority must be set in each of the MEPs monitoring a given Ethernet LSP. How CCM priority is determined is out of scope of this document. Note that the highest priority should be used as the default CCM priority.
- o MEPs must be aware of the reachability parameters of their own and that of the remote MEP. In the case of bidirectional point-to-point PBB-TE connections, this requires that the 3-tuples [ESP-MAC A, ESP-MAC B, ESP-VID1] and [ESP-MAC B, ESP-MAC A, ESP-VID2] are configured in each MEP, where the ESP-MAC A is the same as the local MEP's MAC address and ESP-MAC B is the same as remote MEP's MAC address. The GMPLS Ethernet Label format, as defined in [RFC6060], consists of the ESP-MAC DA and ESP-VID. Hence the necessary reachability parameters for the MEPs can be obtained from the Ethernet Labels (i.e., carried in the downstream and upstream labels). In the case of point-to-point VLAN connections, MEPs need to be provisioned with the VLAN identifiers only, which can be derived similarly from the Ethernet Labels.

Based on the procedures described in [RFC6060] for bidirectional PBB-TE Ethernet LSP establishment, the Ethernet OAM configuration procedures are as follows:

When the RSVP-TE signaling is initiated for the bidirectional Ethernet LSP the local node generates a Path message and:

- o Allocates an Upstream Label formed by combining its MAC address (ESP-MAC A) and locally selected VID (ESP-VID1), which will be used to receive traffic;
- o MUST include the OAM Configuration TLV with OAM Type set to Ethernet OAM in the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRIBUTES Object;
- o MUST include the OAM Function Flags Sub-TLV in the OAM Configuration TLV and set the OAM function flags as needed;
- o MUST include an Ethernet OAM Configuration Sub-TLV in the OAM Configuration TLV that specifies the CCM Interval and MD Level;

- o MAY add an MD Name Sub-TLV (optional) and MUST add a Short MA Name Sub-TLV (required) to the Ethernet OAM Configuration Sub-TLV, that will unambiguously identify a Maintenance Association for this specific PBB-TE connection. Note that values for these parameters may be derived from the GMPLS LSP identification parameters;
- o MUST include a MEP ID Sub-TLV in the Ethernet OAM Configuration Sub-TLV and select two distinct integer values to identify the local and remote MEPs within the Maintenance Association created for monitoring of the point-to-point PBB-TE connection.

Once the remote node receives the Path message, it can use the UPSTREAM_LABEL to extract the reachability information of the initiator. Then it allocates a Label by selecting a local MAC address (ESP-MAC B) and VID (ESP-VID2) that will be used to receive traffic. These parameters determine the reachability information of the local MEP. That is, the 3-tuples [ESP-MAC A, ESP-MAC B, ESP-VID1] and [ESP-MAC B, ESP-MAC A, ESP-VID2] are derived from the Ethernet Labels. In addition, the information received in the Ethernet OAM Configuration TLV is used to configure the local MEP.

Once the Resv message successfully arrives to the initiator, this end can extract the remote side's reachability information from the Label Object and therefore it has all the information needed to properly configure its local MEP.

3.2. OAM Configuration TLV

This TLV is specified in [RFC7260] and is used to select which OAM technology/method should be used for the LSP. In this document, a new OAM Type: Ethernet OAM is defined. IANA is requested to allocate OAM Type 1 for Ethernet OAM in the RSVP-TE OAM Configuration Registry.

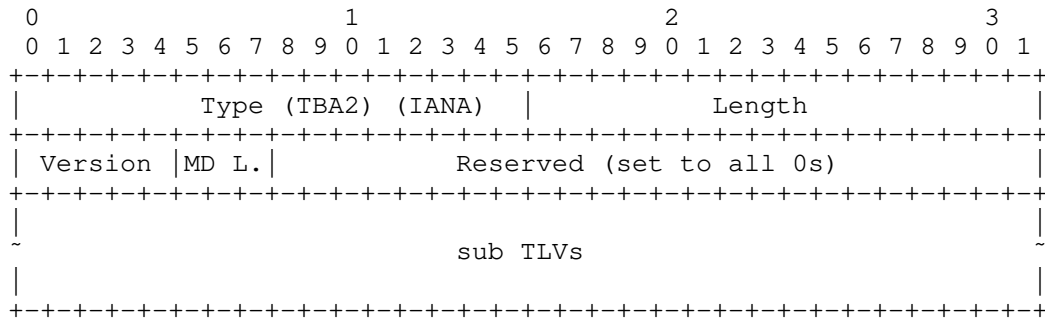
RSVP-TE OAM Configuration Registry

OAM Type	Description
-----	-----
TBA1	Ethernet OAM

The receiving node, when the Ethernet OAM Type is requested, should look for the corresponding technology specific Ethernet OAM Configuration Sub-TLV.

3.3. Ethernet OAM Configuration Sub-TLV

The Ethernet OAM Configuration Sub-TLV (depicted below) is defined for Ethernet OAM specific configuration parameters. The Ethernet OAM Configuration Sub-TLV, when used, MUST be carried in the OAM Configuration TLV. This new sub-TLV accommodates Ethernet OAM information and carries sub-TLVs.



Type: indicates a new type: the Ethernet OAM Configuration Sub-TLV. IANA is requested to assign a value from the "Sub-TLV" space in the "RSVP-TE OAM Configuration Registry".

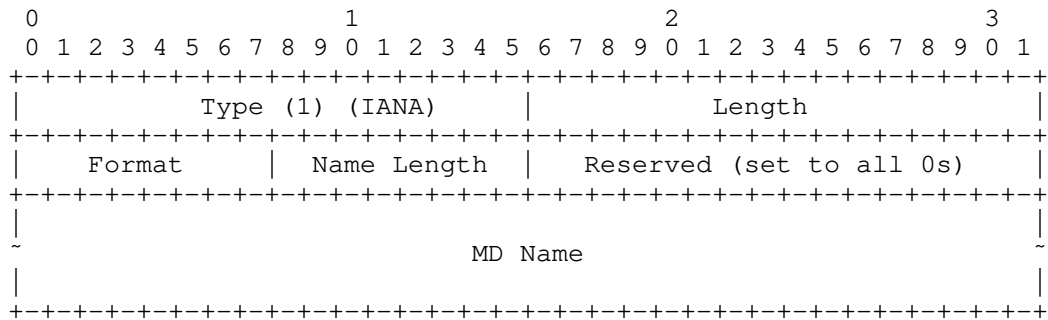
Length: indicates the total length of the TLV including padding and including the Type and Length fields.

Version: identifies the CFM protocol version according to [IEEE.802.1Q-2011]. If a node does not support a specific CFM version an error MUST be generated: "OAM Problem/Unsupported OAM Version"

MD L. (MD Level): indicates the desired MD Level. Possible values are defined according to [IEEE.802.1Q-2011]. If a node does not support a specific MD Level an error MUST be generated: "OAM Problem/Unsupported MD Level".

3.3.1. MD Name Sub-TLV

The optional MD Name Sub-TLV is depicted below, it MAY be used for MD naming.



Type: 1, MD Name Sub-TLV. IANA is requested to maintain an Ethernet TLV Type space in the "RSVP-TE OAM Configuration Registry" for the sub-TLV types carried in the Ethernet OAM Configuration Sub-TLV.

Length: indicates the total length of the TLV including padding and including the Type and Length fields.

Format: according to [IEEE.802.1Q-2011].


Name Length: the length of the MD Name field in bytes. This is necessary to allow non 4 byte padded MD Name lengths.

MD Name: variable length field, formatted according to the format specified in the Format field.

If an undefined Format is specified an error MUST be generated: "OAM Problem/Unknown MD Name Format". Also the combined length of MD Name and Short MA Name MUST be less or equal to 44bytes, if this is violated an error MUST be generated: "OAM Problem/Name Length Problem". Note it is allowed to have no MD Name, therefore the MD Name Sub-TLV is optional. In this case the MA Name must uniquely identify a Maintenance Association.

3.3.2. Short MA Name Sub-TLV

The Short MA Name Sub-TLV is depicted below. This sub-TLV MUST be present in the Ethernet OAM Configuration Sub-TLV.



Type: 2, Short MA Name Sub-TLV. IANA is requested to maintain an Ethernet TLV Type space in the "RSVP-TE OAM Configuration Registry" for the sub-TLV types carried in the Ethernet OAM Configuration Sub-TLV.

Length: indicates the total length of the TLV including padding and including the Type and Length fields.

Format: according to [IEEE.802.1Q-2011].


Name Length: the length of the MA Name field in bytes. This is necessary to allow non 4 byte padded MA Name lengths.

Short MA Name: variable length field formatted according to the format specified in the Format field.

If an undefined Format is specified an error MUST be generated: "OAM Problem/Unknown MA Name Format". Also the combined length of MD Name and Short MA Name MUST be less or equal to 44bytes, if this is violated an error MUST be generated: "OAM Problem/Name Length Problem". Note it is allowed to have no MD Name, in this case the MA Name MUST uniquely identify a Maintenance Association.

3.3.3. MEP ID Sub-TLV

The MEP ID Sub-TLV is depicted below. This sub-TLV MUST be present in the Ethernet OAM Configuration Sub-TLV.



Type: 3, MEP ID Sub-TLV. IANA is requested to maintain an Ethernet TLV Type space in the "RSVP-TE OAM Configuration Registry" for the sub-TLV types carried in the Ethernet OAM Configuration Sub-TLV.

Length: indicates the total length of the TLV including padding and including the Type and Length fields.


Local MEP ID: a 16 bit integer value in the range 1-8191 of the MEP ID on the initiator side.

Remote MEP ID: a 16 bit integer value in the range 1-8191 of the MEP ID to be set for the MEP established at the receiving side. This value is determined by the initiator node. This is possible, since a new MAID is assigned to each PBB-TE connection, and MEP IDs must be only unique within the scope of the MAID.

Two flags are defined Transmit (T) and Receive (R). When T is set the corresponding MEP MUST send OAM packets. When R is set the corresponding MEP MUST expect to receive OAM packets. These flags are used to configure the role of MEPs.

3.3.4. Continuity Check (CC) Sub-TLV

The Continuity Check (CC) Sub-TLV is depicted below. This sub-TLV MUST be present in the Ethernet OAM Configuration Sub-TLV.



Type: 4, Continuity Check (CC) Sub-TLV. IANA is requested to maintain an Ethernet TLV Type space in the "RSVP-TE OAM Configuration Registry" for the sub-TLV types carried in the Ethernet OAM Configuration Sub-TLV.

Length: indicates the total length of the TLV including padding and including the Type and Length fields.

Prio: Indicates the priority to be set for CCM frames. In Ethernet, 3 bits carried in VLAN TAGs identify priority information. Setting the priority is optional. If the most significant bit is set to zero, the subsequent 3 priority bits will be ignored, and priority bits of the Ethernet CCM frame will be set based on default values specified in the Ethernet nodes. If the most significant bit is set to 1, the subsequent 3 bits will be used to set the priority bits of the Ethernet CCM frame.

CCM I (CCM Interval): CCM Interval, it MUST be set according to the 3 bit encoding [IEEE.802.1Q-2011] shown in Table 1. As a consequence the most significant bit will be set to 0. Four bits are allocated to support the configuration of CCM intervals that may be specified in the future. If a node does not support the requested CCM Interval an error MUST be generated: "OAM Problem/Unsupported CC Interval".

3.4. Pro-active Performance Monitoring

Ethernet OAM functions for Performance Monitoring (PM) allow measurements of different performance parameters including Frame Loss Ratio, Frame Delay and Frame Delay variation as defined in [ITU-T.Y.1731-2011]. Only a subset of PM functions are operated in a pro-active fashion to monitor the performance of the connection continuously. Pro-active PM supports Fault Management functions, by providing an indication of decreased service performance and as such may provide triggers to initiate recovery procedures.

While on demand PM functions are, for the purposes of this document, always initiated by management commands, for pro-active PM, it may be desirable to utilize the control plane for configuration and activation together with Fault Management functions such as the Continuity Check.

[ITU-T.Y.1731-2011] defines dual-ended Loss Measurement as pro-active OAM for performance monitoring and as a PM function applicable to fault management. For dual-ended Loss Measurement each MEP piggy-backs transmitted and received frame counters on CC messages; to support and synchronize bidirectional Loss Measurements at the MEPs. Dual-ended Loss Measurement is supported by setting the Performance Monitoring/Loss OAM Function Flag and the Continuity Check Flag in the OAM Function Flags Sub-TLV [RFC7260], and configuring the Continuity Check functionality by including the Ethernet OAM Configuration Sub-TLV. No additional configuration is required for this type of Loss Measurement.

3.5. Summary of Ethernet OAM configuration errors

In addition to error values specified in [RFC7260] this document defines the following values for the "OAM Problem" Error Code.

- o If a node does not support a specific CFM version, an error MUST be generated: "OAM Problem/Unsupported OAM Version".
- o If a node does not support a specific MD Level, an error MUST be generated: "OAM Problem/Unsupported MD Level".
- o If an undefined MD name format is specified, an error MUST be generated: "OAM Problem/Unknown MD Name Format".
- o If an undefined MA name format is specified, an error MUST be generated: "OAM Problem/Unknown MA Name Format".
- o The combined length of MD Name and Short MA Name must be less or equal to 44bytes, if this is violated an error MUST be generated: "OAM Problem/Name Length Problem".
- o If a node does not support the requested CCM Interval, an error MUST be generated: "OAM Problem/Unsupported CC Interval".

4. IANA Considerations

4.1. RSVP-TE OAM Configuration Registry

IANA maintains the "RSVP-TE OAM Configuration Registry". IANA is requested to assign an "OAM Type" from this registry as follows. Allocate the value TBA1 for "Ethernet OAM" from the "OAM Type Sub-Registry" of the "RSVP-TE OAM Configuration Registry". Allocate type TBA2 for the "Ethernet OAM Configuration Sub-TLV" from the technology-specific range of the "OAM Sub-TLVs Sub-Registry" of the "RSVP-TE OAM Configuration Registry".

RSVP-TE OAM Configuration Registry

OAM Types Sub-Registry

OAM Type Number	Description	Reference
TBA1	Ethernet OAM	[This.ID]

OAM Sub-TLVs Sub-Registry

Sub-TLV Type	Description	Ref.
TBA2	Ethernet OAM Configuration Sub-TLV	[This.ID]

The value of 1 is suggested for TBA1. The value of 32 is suggested for TBA2.

4.2. Ethernet Sub-TLVs Sub-Registry

IANA is requested to maintain an Ethernet Sub-TLVs Sub-Registry in the "RSVP-TE OAM Configuration Registry" for the sub-TLV types carried in the Ethernet OAM Configuration Sub-TLV. This document defines the following types.

Ethernet Sub-TLVs Sub-Registry

Range	Registration Procedures
0-65534	IETF Review
65535-65536	Experimental

Sub-TLV Type	Description	Ref.
0	Reserved	[This.ID]
1	MD Name Sub-TLV	[This.ID]
2	Short MA Name Sub-TLV	[This.ID]
3	MEP ID Sub-TLV	[This.ID]
4	Continuity Check Sub-TLV	[This.ID]
5-65536	Unassigned	[This.ID]

4.3. RSVP Error Code

IANA maintains an Error Code, "OAM Problem" in the "Error Codes and Globally-Defined Error Value Sub-Codes" sub-registry of the "Resource Reservation Protocol (RSVP) Parameters" registry. [RFC7260] defines a set of Error Value sub-codes for the "OAM Problem" Error Code.

This document defines additional Error Values sub-codes for the "OAM Problem" Error Code as summarized below.

Value	Description	Reference
TBA	Unsupported OAM Version	[This.ID]
TBA	Unsupported MD Level	[This.ID]
TBA	Unknown MD Name Format	[This.ID]
TBA	Unknown MA Name Format	[This.ID]
TBA	Name Length Problem	[This.ID]
TBA	Unsupported CC Interval	[This.ID]

5. Security Considerations

This document does not introduce any additional security issue to those discussed in [RFC7260] and [RFC6060].

The signaling of OAM-related parameters and the automatic establishment of OAM entities based on RSVP-TE messages add a new aspect to the security considerations discussed in [RFC3473]. In particular, a network element could be overloaded if a remote attacker targeted that element by sending frequent periodic messages requesting liveliness monitoring of a high number of LSPs. Such an attack can efficiently be prevented when mechanisms for message integrity and node authentication are deployed. Since the OAM configuration extensions rely on the hop-by-hop exchange of exiting RSVP-TE messages, procedures specified for RSVP message security in [RFC2747] can be used to mitigate possible attacks.

For a more comprehensive discussion of GMPLS security and attack mitigation techniques, please see the Security Framework for MPLS and GMPLS Networks [RFC5920].

6. Acknowledgements

The authors would like to thank Francesco Fondelli, Adrian Farrel, Loa Andersson, Eric Gray and Dimitri Papadimitriou for their useful comments.

7. Contributors

- Don Fedyk, don.fedyk@hp.com
- Dinesh Mohan, dinmohan@hotmail.com

8. References

8.1. Normative References

- [IEEE.802.1Q-2011]
IEEE, "IEEE Standard for Local and metropolitan area networks -- Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q, 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC6060] Fedyk, D., Shah, H., Bitar, N., and A. Takacs, "Generalized Multiprotocol Label Switching (GMPLS) Control of Ethernet Provider Backbone Traffic Engineering (PBB-TE)", RFC 6060, March 2011.
- [RFC7260] Takacs, A., Fedyk, D., and J. He, "GMPLS RSVP-TE Extensions for Operations, Administration, and Maintenance (OAM) Configuration", RFC 7260, June 2014.

8.2. Informative References

- [ITU-T.Y.1731-2011]
ITU, "ITU-T Recommendation Y.1731: OAM functions and mechanisms for Ethernet based networks", ITU-T Recommendation Y.1731, 2011.
- [RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.
- [RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC5828] Fedyk, D., Berger, L., and L. Andersson, "Generalized Multiprotocol Label Switching (GMPLS) Ethernet Label Switching Architecture and Framework", RFC 5828, March 2010.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Authors' Addresses

Attila Takacs
Ericsson
Konyves Kalman krt. 11.
Budapest 1097
Hungary

Email: attila.takacs@ericsson.com

Balazs Peter Gero
Ericsson
Konyves Kalman krt. 11.
Budapest 1097
Hungary

Email: balazs.peter.gero@ericsson.com

Hao Long
Huawei
PR China

Email: lonho@huawei.com