Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 7, 2009

A. Takacs
B. Gero
Ericsson
D. Fedyk
D. Mohan
Nortel

November 3, 2008

GMPLS RSVP-TE Extensions for Ethernet OAM Configuration
draft-takacs-ccamp-rsvp-te-eth-oam-ext-03

Status of this Memo

Abstract

The GMPLS controlled Ethernet Label Switching (GELS) work is extending 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 to support the setup of the associated Ethernet OAM (CFM and Y.1731) entities adding a technology specific TLV to [OAM-CONF-FWK].

Changes from previous version

Technology independent extensions were moved to a separate framework document leaving only the Ethernet specific extensions in this document.

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

Table of Contents

1. Background

Provider Backbone Bridging - Traffic Engineering (PBB-TE) [IEEE-PBBTE] decouples the Ethernet data and control planes by explicitly supporting external control/management mechanisms to configure static filtering entries in bridges and create explicitly routed Ethernet connections. In addition PBB-TE defines mechanisms for 1:1 protection switching of bidirectional Ethernet connections.

Ethernet Connectivity Fault Management (CFM) defines an adjunct connectivity monitoring OAM flow to check the liveliness of Ethernet networks [IEEE-CFM]. With PBB-TE Ethernet networks will support explicitly-routed Ethernet connections. CFM can be used to track the liveliness of PBB-TE connections and detect data plane failures.

In IETF the GMPLS controlled Ethernet Label Switching (GELS) [GELS-Framework] work is extending the GMPLS control plane to support the establishment of point-to-point PBB-TE data plane connections. We refer to GMPLS established PBB-TE connections as Ethernet LSPs. GELS enables the application of MPLS-TE and GMPLS provisioning and recovery features in Ethernet networks.

2. Overview of Ethernet OAM operation in PBB-TE networks

For the purposes of this document, we only discuss Ethernet OAM [IEEE-CFM] aspects that are relevant for the connectivity monitoring of point-to-point PBB-TE connections.

PBB-TE [IEEE-PBBTE] 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.

Note that although it would be possible to use GMPLS to setup a single unidirectional ESP, the Ethernet OAM mechanisms are only full functional when bidirectional connections are established with co-routed ESPs. Hence, we focus on bidirectional point-to-point PBB-TE connections.

At both ends of the bidiretional 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.

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

#	CCM Interval (CCI)	3 bit encoding
0	Invalid	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 not received by one of the MEPs it declares a connectivity failure and signals the failure in subsequent CCM messages, by setting the Remote Defect Indicator (RDI) bit, to the remote MEP. If a MEP receives a CCM message with RDI set it immediately declares failure. The detection of a failure may trigger protection switching mechanisms or may be signalled to a management system. However, what happens once a failure is detected is out of the scope of this document.

3. GMPLS RSVP-TE Extensions

3.1. Operation overview

To simplify the configuration of connectivity monitoring, when an Ethernet LSP is signalled 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.

- 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 formating rules for these names have been defined by [IEEE-CFM]. Since these information is also carried in all CCM messages, the combined length of the Names is limited to 44 bytes. 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-CFM] allows differentiation of 8 levels. How the value of the MD Level is determined is out of scope of this document. Note that most probably for all Ethernet LSPs a single (default) MD Level will be used.

- 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 CCM priority to be set by MEPs for the CCM frames can 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.

- o MEPs must be aware of their own and the reachability parameters 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 and ESP-MAC B is the same as remote MEP's MAC. The GMPLS Ethernet Label for forwarding, as defined in [GELS-PBBTE], consists of the ESP-MAC DA and ESP-VID. Hence the necessary reachability parameters for the MEPs can be obtained from Ethernet Labels (i.e., carried in the "downstream" and upstream labels).

Assuming the procedures described in [GELS-PBBTE] for bidirectional Ethernet LSP establishment the MEP configuration should be as follows. When the RSVP-TE signalling is initiated for the bidirectional Ethernet LSP the local node generates a Path message and:

- Allocates an Upstream Label from its MAC address (ESP-MAC A) and locally selected VID (ESP-VID1), that it would like to use to receive traffic;

- Inserts an Ethernet OAM Configuration TLV in the LSP_ATTRIBUTES object, specifying the CCM Interval and MD Level;

- Adds an MD Name Sub-TLV (optional) and a Short MA Name Sub-TLV to the Ethernet OAM Configuration 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;

- Adds a MEP ID Sub-TLV to the Ethernet OAM Configuration TLV. It selects 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 the MAC address (ESP-MAC B) and VID (ESP-VID2) it would like to use 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 it can extract the remote side's reachability information from the LABEL object whereby this node has also obtained all the information needed to establish its local MEP.

Once the MEPs are established the monitoring of the LSP is operational. In certain situations, e.g., maintenance, re-

optimisation of LSPs, it is desirable to explicitly enable or disable the monitoring of LSPs (i.e., start/stop exchanging CC messages). To allow administrative control of LSP monitoring the "Monitoring Disabled" (M) bit in the ADMIN_STATUS Object is used [OAM-CONF-FWK].

3.2. OAM Configuration TLV

This TLV is specified in [OAM-CONF-FWK] 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.

OAM Type	Description
0	Reserved
1	Ethernet OAM
2-256	Reserved

The receiving node when the Ethernet OAM Type is requested should look for the corresponding technology specific Ethernet OAM configuration TLV.

3.3. Ethernet OAM Configuration TLV

The Ethernet OAM Configuration TLV (depicted below) is defined for Ethernet OAM specific configuration parameters. The Ethernet OAM Configuration TLV is carried in the LSP_ATTRIBUTES object in Path messages. This new TLV accommodates generic Ethernet OAM information and carries sub-TLVs.

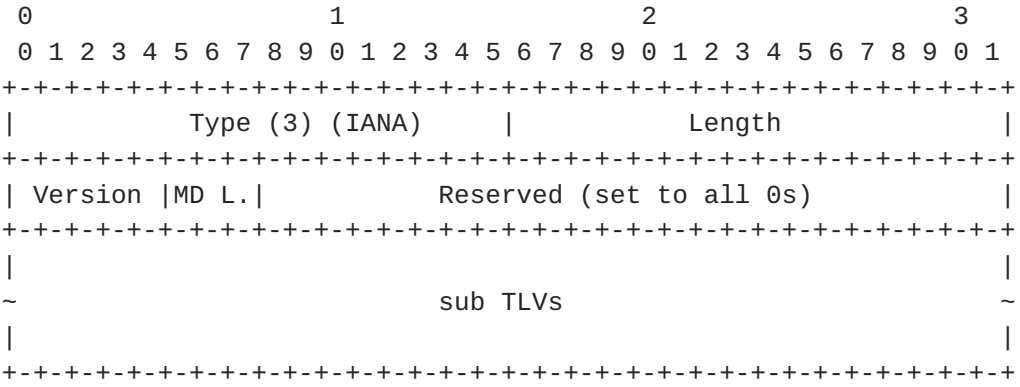

Type: indicates a new type: the Ethernet OAM Configuration TLV (3) (IANA to define).

Length: indicates the total length including sub-TLVs.

Version: identifies the CFM protocol version according to [IEEE-CFM]. 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. The values are according to [IEEE-CFM]. If a node does not support a specific MD Level an error must be generated: "OAM Problem/Unsupported OAM Level".

3.3.1. MD Name Sub-TLV

The optional MD Name sub-TLV is depicted below.

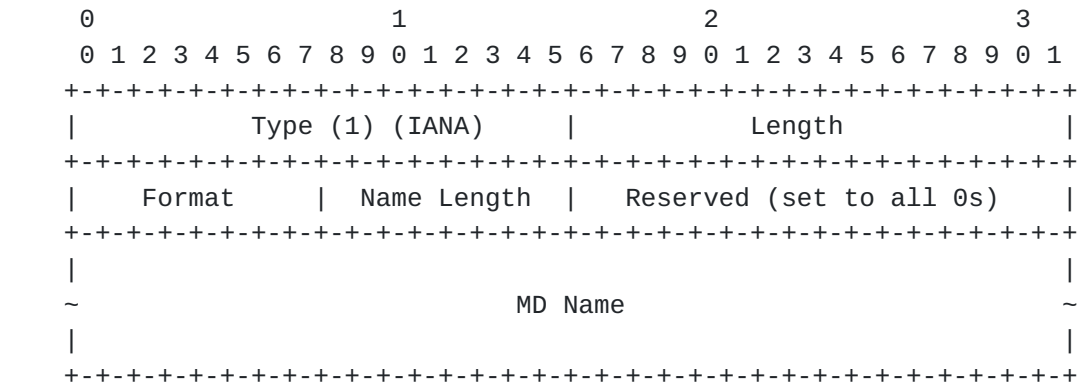

Type: 1, MD Name Sub-TLV (IANA).

Length: indicates the total length of the TLV including padding.

Format: according to [IEEE-CFM].

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 that it is allowed to have no MD Name, as such 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.

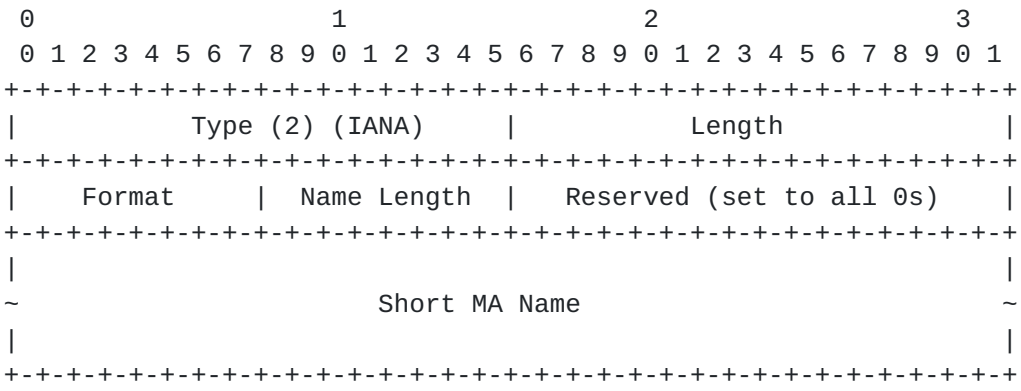

Type: 2, Short MA Name Sub-TLV (IANA)

Length: indicates the total length of the TLV including padding.

Format: according to [IEEE-CFM].

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 that 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.

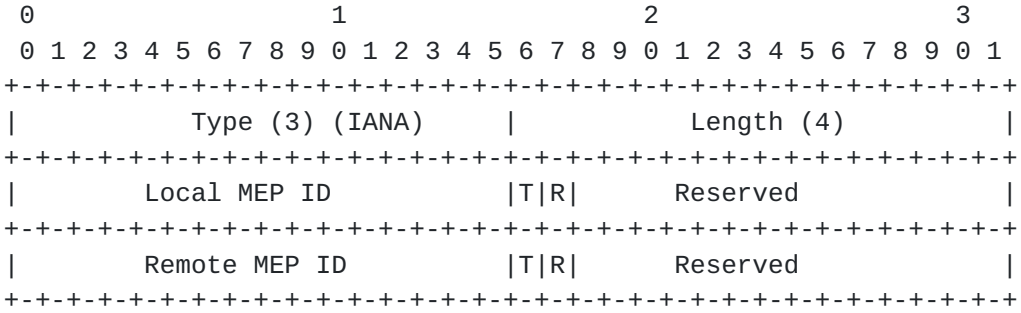

Type: 3, MEP ID Sub-TLV (IANA)

Length: indicates the total length of the TLV including padding.

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.

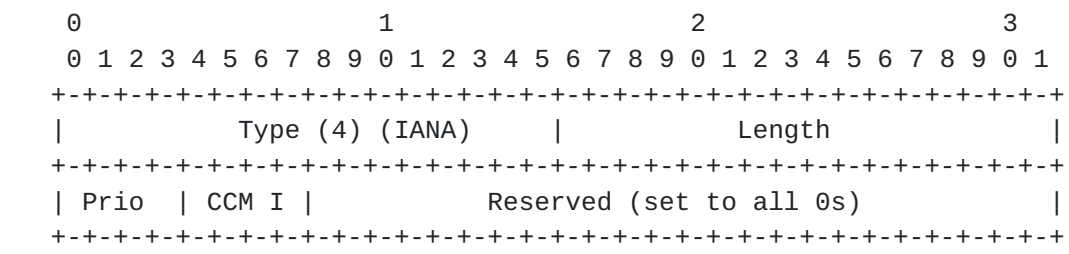

Prio: Indicates the priority to be set for CCM frames. In Ethernet 3 bits carried in VLAN TAGs identify priority information.

CCM I (CCM Interval): CCM Interval, according to the 3 bit encoding [IEEE-CFM] shown in Table 1. If a node does not support the requested CCM Interval an error must be generated: "OAM Problem/ Unsupported CC Interval".

3.4. Ethernet OAM configuration errors

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

- If a node does not support a specific CFM version an error must be generated: "OAM Problem/Unsupported OAM Version".

- If a node does not support a specific MD Level an error must be generated: "OAM Problem/Unsupported OAM Level".

- If an undefined MD name format is specified an error must be generated: "OAM Problem/Unknown MD Name Format".

- If an undefined MA name format is specified an error must be generated: "OAM Problem/Unknown MA Name Format".

- If 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".

- If a node does not support the requested CCM Interval an error must be generated: "OAM Problem/Unsupported CC Interval".

4. IANA Considerations

This document specifies a new Ethernet OAM Configuration TLV to be carried in the OAM Configuration TLV in LSP_ATTRIBUTES and LSP_REQUIRED_ATTRIBUTES objects in Path messages.

The following values need to be assigned under the Error Code: "OAM Problem": "Unsupported OAM Version", "Unsupported OAM Level", "Unknown MD Name Format", "Unknown MA Name Format", "Name Length Problem", "Unsupported CC Interval".

5. Security Considerations

The signalling of OAM related parameters and the automatic establishment of OAM entities introduces additional security considerations to those discussed in [RFC3473]. In particular, a network element could be overloaded, if an attacker would request liveliness monitoring, with frequent periodic messages, for a high number of LSPs, targeting a single network element.

Security aspects will be covered in more detailed in subsequent versions of this document.

6. Acknowledgements

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

7. References

[GELS-Framework]
"GMPLS Ethernet Label Switching Architecture and Framework", Internet Draft, work in progress.

[GELS-PBBTE]
"GMPLS control of Ethernet PBB-TE", Internet Draft, work in progress.

[GMPLS-OAM]
"OAM Requirements for Generalized Multi-Protocol Label Switching (GMPLS) Networks", Internet Draft, work in progress.

[IEEE-CFM]
"IEEE 802.1ag, Draft Standard for Connectivity Fault Management", work in progress.

[IEEE-PBBTE]
"IEEE 802.1Qay Draft Standard for Provider Backbone Bridging Traffic Engineering", work in progress.

[OAM-CONF-FWK]
"OAM Configuration Framework for GMPLS RSVP-TE", Internet Draft, work in progress.

[RFC3469] "Framework for Multi-Protocol Label Switching (MPLS)-based Recovery", RFC 3469, February 2003.

[RFC3471] "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

[RFC4377] "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.

[RFC4420] "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4420, February 2006.

Authors' Addresses

Attila Takacs
Ericsson
Laborc u. 1.
Budapest, 1037
Hungary

Email: attila.takacs@ericsson.com

Balazs Gero
Ericsson
Laborc u. 1.
Budapest, 1037
Hungary

Email: balazs.gero@ericsson.com

Don Fedyk
Nortel
600 Technology Park Drive
Billerica, MA 01821
USA

Email: dwfedyk@nortel.com

Dinesh Mohan
Nortel
3500 Carling Ave
Ottawa, ON, K2H8E9
Canada

Email: mohand@nortel.com

Full Copyright Statement

Intellectual Property