Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 14, 2012

A. Takacs
B. Gero
Ericsson
H. Long
Huawei

January 11, 2012

GMPLS RSVP-TE Extensions for Ethernet OAM Configuration
draft-ietf-ccamp-rsvp-te-eth-oam-ext-07

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 to support the setup of the associated Ethernet OAM (CFM and Y.1731) entities defining Ethernet technology specific TLV based on [OAM-CONF-FWK].

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

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 July 14, 2012.

Copyright Notice

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 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], including the monitoring of explicitly-routed Ethernet connections.

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.

2. Overview of Ethernet OAM operation

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

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 only.

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 is identified with a single VLAN forwarding traffic in both directions or with two VLANs each forwarding traffic in a single direction. Hence instead of the 3-tuples of the PBB-TE case MEPs must be provisioned with the proper VLAN information, otherwise the same MD Level, MAID, MEP ID configuration is required in this case as well.

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	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 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 signaled to a management system. However, what happens once a failure is detected is out of the scope of this document.

At each transit node Maintenance Intermediate Points (MIPs) can be established to help failure localization by supporting link trace and loop back functions. MIPs need to be provisioned with a subset of 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 Attributes Flags: "OAM MEP entities desired" and "OAM MIP entities desired", described in [OAM-CONF-FWK] are used to signal that the respective OAM entities must be established. Subsequently, an OAM Configuration TLV is added to the LSP_ATTRIBUTES Object specifying that Ethernet OAM is to be setup for the LSP. The below detailed Ethernet OAM specific information is carried in the new Ethernet OAM Configuration sub-TLV.

- 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 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.

- 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.

- 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 within a network domain.

- 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.

- 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. Note that the highest priority is used as the default CCM priority.

- 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 address and ESP-MAC B is the same as remote MEP's MAC address. The GMPLS Ethernet Label for forwarding, 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 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, which can be derived similarly from the Ethernet Label.

Assuming the procedures described in [RFC6060] for bidirectional PBB-TE Ethernet LSP establishment the MEP configuration should be as follows. When the RSVP-TE signaling 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), which will be used to receive traffic;

- Inserts the OAM Configuration TLV with OAM Type set to Ethernet OAM in the LSP_ATTRIBUTES object;

- Adds the OAM Function Flags sub-TLV in the OAM Configuration TLV and sets the OAM function flags as needed;

- Adds an Ethernet OAM Configuration sub-TLV in the OAM Configuration TLV that specifies 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.

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 within the OAM Configuration TLV in the LSP_ATTRIBUTES or LSP_REQUIRED_ATTRUIBUTES object in Path messages. This new TLV accommodates generic Ethernet OAM information and carries sub-TLVs.

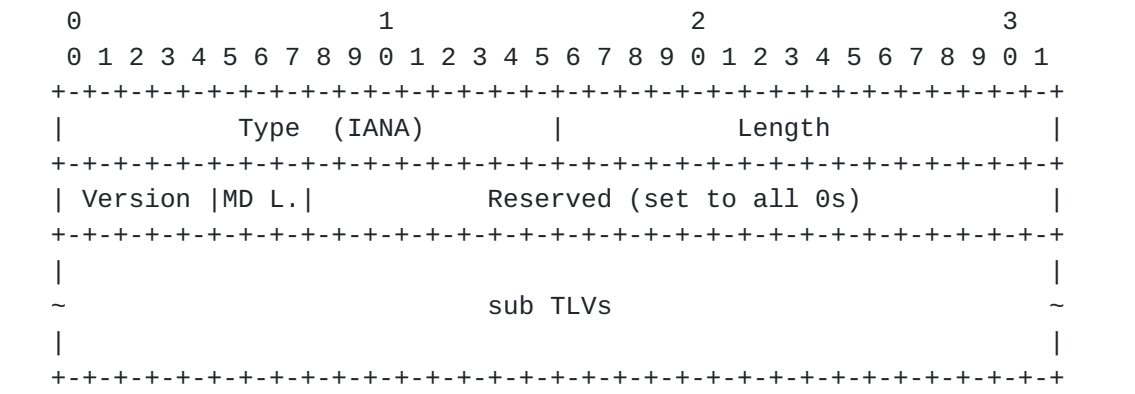

Type: indicates a new type: the Ethernet OAM Configuration TLV. IANA is requested to assigne a value from the "OAM Type sub-TLV" space in the "RSVP-TE OAM Configuration Registry".

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.

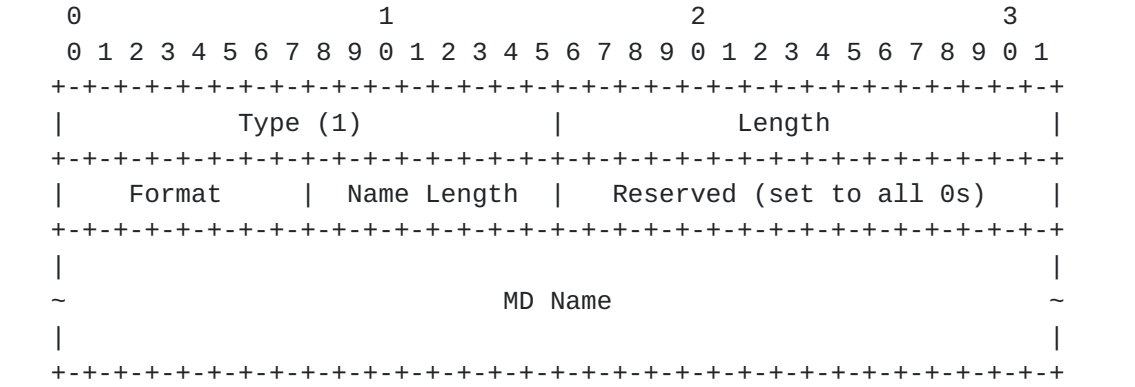

Type: 1, MD Name Sub-TLV.

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.

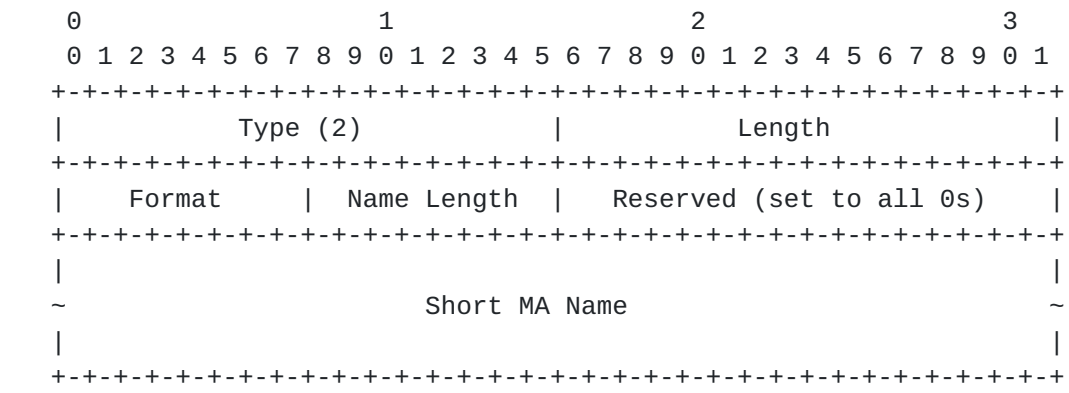

Type: 2, Short MA Name Sub-TLV.

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.

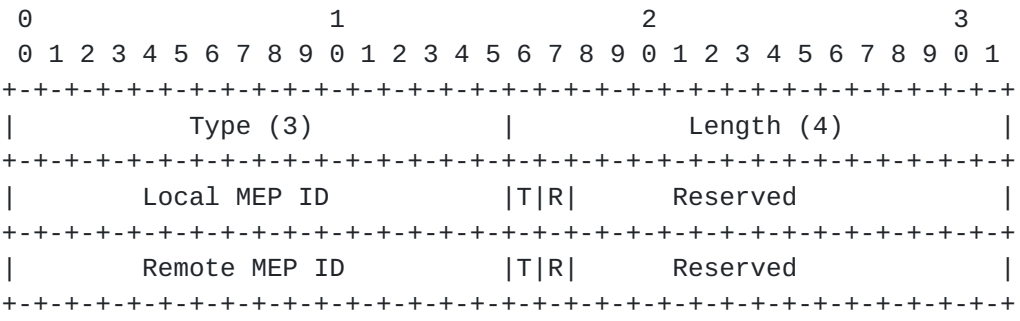

Type: 3, MEP ID Sub-TLV.

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.

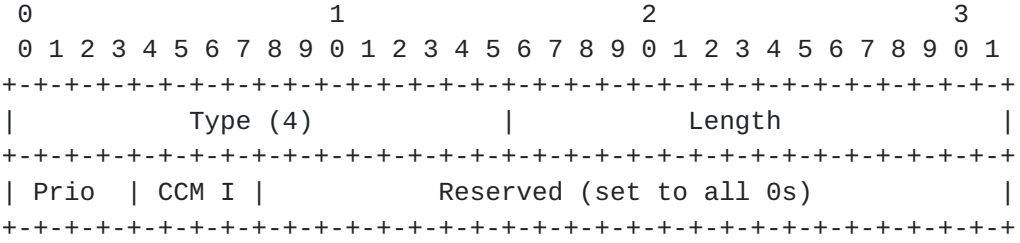

Type: 4, Continuity Check (CC) sub-TLV.

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. 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 the ITU-T Y.1731 recommendation. 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 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 Continuity Check.

ITU-T Y.1731 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 invoked by setting the Performance Monitoring/Loss OAM Function Flag in the OAM Function Flags Sub-TLV [OAM-CONF-FWK]. Besides configuring the Continuity Check functionality, no additional configuration is required for this type of Loss Measurement.

3.5. 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 the Ethernet OAM Configuration sub-TLV to be carried in the OAM Configuration TLV in LSP_ATTRIBUTES and LSP_REQUIRED_ATTRIBUTES objects in Path messages.

IANA is requested to allocate the value 1 for Ethernet OAM from the OAM Type space in the "RSVP-TE OAM Configuration Registry" and allocate type 1 for the Ethernet OAM Configuration sub-TLV from the OAM Type sub-TLV space in the "RSVP-TE OAM Configuration Registry".

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

This document does not introduce any additional security issuse to those discussed in [OAM-CONF-FWK].

6. Acknowledgements

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

Contributors

Don Fedyk
Alcatel-Lucent
Groton, MA 01450
USA
Email: donald.fedyk@alcatel-lucent.com

Dinesh Mohan

7. References

7.1. Normative References

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

[RFC5828] "GMPLS Ethernet Label Switching Architecture and Framework", RFC 5828, March 2010.

[RFC6060] "Generalized Multiprotocol Label Switching (GMPLS) Control of Ethernet Provider Backbone Traffic Engineering (PBB-TE)", RFC 6060.

7.2. Informative References

[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.

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

Hao Long
Huawei

Email: lonho@huawei.com