

Internet Engineering Task Force (IETF)
Request for Comments: 6060
Category: Standards Track
ISSN: 2070-1721

D. Fedyk
Alcatel-Lucent
H. Shah
Ciena
N. Bitar
Verizon
A. Takacs
Ericsson
March 2011

Generalized Multiprotocol Label Switching (GMPLS) Control of
Ethernet Provider Backbone Traffic Engineering (PBB-TE)

Abstract

This specification is complementary to the GMPLS Ethernet Label Switching Architecture and Framework and describes the technology-specific aspects of GMPLS control for Provider Backbone Bridge Traffic Engineering (PBB-TE). The necessary GMPLS extensions and mechanisms are described to establish Ethernet PBB-TE point-to-point (P2P) and point-to-multipoint (P2MP) connections. This document supports, but does not modify, the standard IEEE data plane.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6060>.

Copyright Notice

Copyright (c) 2011 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
1.1. Co-Authors	3
2. Terminology	4
2.1. PBB-TE and GMPLS Terminology	5
2.2. Conventions Used in This Document	6
3. Creation and Maintenance of PBB-TE Paths Using GMPLS	6
3.1. Shared Forwarding	9
3.2. P2P Connections Procedures for Shared Forwarding	10
4. Specific Procedures	10
4.1. P2P Ethernet LSPs	10
4.1.1. P2P Path Maintenance	11
4.2. P2MP Ethernet-LSPs	12
4.3. PBB-TE Ethernet Label	12
4.4. Protection Paths	13
4.5. Service Instance Identification	13
5. Error Conditions	15
5.1. ESP-VID-Related Errors	15
5.1.1. Invalid ESP-VID Value in the PBB-TE Ethernet Label	15
5.1.2. Allocated ESP-VID Range is Exhausted	16
5.2. Invalid MAC Address	16
6. Security Considerations	16
7. IANA Considerations	17
8. References	17
8.1. Normative References	17
8.2. Informative References	19
9. Acknowledgments	19

1. Introduction

The IEEE 802.1 Provider Backbone Bridge Traffic Engineering (PBB-TE) [IEEE802.1Qay] standard supports the establishment of explicitly routed traffic engineered paths within Provider Backbone Bridged (PBB) networks. PBB-TE allows the disabling of:

- the Spanning Tree Protocol
- unknown destination address forwarding
- source address learning

for administratively selected VLAN Identifiers. With PBB-TE an external provisioning system or control plane can be used to configure static entries in the managed objects of bridges and so establish traffic engineered paths in the network.

Generalized MPLS (GMPLS) [RFC3945] is a family of control plane protocols designed to operate in connection oriented and traffic engineering transport networks. GMPLS is applicable to a range of network technologies including L2SC networks (Layer 2 Switching Capable). The purpose of this document is to specify extensions for a GMPLS-based control plane to manage PBB-TE explicitly routed traffic engineered paths. This specification is complementary to the GMPLS Ethernet Label Switching Architecture and Framework document [RFC5828].

1.1. Co-Authors

This document is the result of a large team of authors and contributors. The following is a list of the co-authors:

David Allan
Ericsson
EMail: david.i.allan@ericsson.com

Diego Caviglia
Ericsson
Via Negrone 1/A
Genoa, Italy 16153
EMail: diego.caviglia@ericsson.com

Alan McGuire
BT Group PLC
OP6 Polaris House,
Austral Park, Martlesham Heath,
Ipswich, Suffolk, IP5 3RE, UK
EMail: alan.mcguire@bt.com

Nurit Sprecher
Nokia Siemens Networks,
GmbH & Co. KG
COO RTP IE Fixed
3 Hanagar St. Neve Ne'eman B,
45241 Hod Hasharon, Israel
EMail: nurit.sprecher@nsn.com

Lou Berger
LabN Consulting, L.L.C.
Phone: +1-301-468-9228
EMail: lberger@labn.net

2. Terminology

In addition to well-understood GMPLS terms, this memo uses the following terminology from IEEE 802.1 [IEEE802.1ah] [IEEE802.1Qay]:

- BCB	Backbone Core Bridge
- BEB	Backbone Edge Bridge
- B-MAC	Backbone MAC
- B-VID	Backbone VLAN ID
- B-VLAN	Backbone VLAN
- CBP	Customer Backbone Port
- CCM	Continuity Check Message
- CNP	Customer Network Port
- C-MAC	Customer MAC
- C-VID	Customer VLAN ID
- C-VLAN	Customer VLAN
- ESP	Ethernet Switched Path
- ESP-MAC SA	ESP Source MAC Address
- ESP-MAC DA	ESP Destination MAC Address
- ESP-VID	ESP VLAN ID
- Eth-LSP	Ethernet Label Switched Path
- IB-BEB	A BEB comprised of both I- and B-components
- I-SID	Ethernet Service Instance Identifier
- TAG	An Ethernet Header Field with Type and Values
- MAC	Media Access Control
- PBB	Provider Backbone Bridges
- PBB-TE	Provider Backbone Bridges Traffic Engineering
- PIP	Provider Instance Port
- PNP	Provider Network Port
- PS	Protection Switching
- P2P	Point-to-Point
- P2MP	Point-to-Multipoint
- SVL	Shared VLAN Learning

- TESI Traffic Engineering Service Instance
- VID VLAN ID
- VIP Virtual Instance Port
- VLAN Virtual LAN

2.1. PBB-TE and GMPLS Terminology

The PBB-TE specification [IEEE802.1Qay] defines some additional terminology to clarify the PBB-TE functions. We repeat these here in expanded context to translate from IEEE to GMPLS terminology. The terms "bridge" and "switch" are used interchangeably in this document. The signaling extensions described here apply equally well to a PBB-TE-capable bridge supporting GMPLS signaling or to a GMPLS-capable switch supporting Ethernet PBB-TE forwarding.

- Ethernet Switched Path (ESP):

A provisioned traffic engineered unidirectional connectivity path between two or more Customer Backbone Ports (CBPs) that extends over a Provider Backbone Bridge Network (PBBN). The path is identified by the 3-tuple <ESP-MAC DA, ESP-MAC SA, ESP-VID>. An ESP is point-to-point (P2P) or point-to-multipoint (P2MP). An ESP is analogous to a (unidirectional) point-to-point or point-to-multipoint LSP. We use the term Ethernet-LSP (Eth-LSP) for GMPLS established ESPs.

- Point-to-Point ESP:

An ESP between two CBPs. The ESP-DA and the ESP-SA in the ESP's 3-tuple identifier are the individual MAC addresses of the two CBPs.

- Point-to-Multipoint ESP:

An ESP among one root CBP and n leaf CBPs. The ESP-DA in the ESP's 3-tuple identifier is a group MAC address identifying the n leaf CBPs, and the ESP-SA is the individual MAC address of the root.

- Point-to-Point PBB-TE Service Instance (P2P TESI):

A service instance supported by two point-to-point ESPs where the ESPs' endpoints have the same CBP MAC addresses. The two unidirectional ESPs are forming a bidirectional service. The PBB-TE standard [IEEE802.1Qay] notes the following: for reasons relating to TE service monitoring diagnostics, operational simplicity, etc., the IEEE PBB-TE standard assumes that the point-to-point ESPs associated with a point-to-point TESI are

co-routed. Support for a point-to-point TE services that comprises non-co-routed ESPs is problematic, and is not defined in this standard. Hence, a GMPLS bidirectional LSP is analogous to a P2P TE Service Instance. We use the term "bidirectional Ethernet-LSP" for GMPLS-established P2P PBB-TE Service Instances.

2.2. Conventions Used in This Document

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].

3. Creation and Maintenance of PBB-TE Paths Using GMPLS

IEEE PBB-TE is a connection-oriented Ethernet technology. PBB-TE ESPs are created bridge by bridge (or switch by switch) by simple configuration of Ethernet forwarding entries. This document describes the use of GMPLS as a valid control plane for the setup, teardown, protection, and recovery of ESPs and TESIs and specifies the required RSVP-TE extensions for the control of PBB-TE Service Instances.

PBB-TE ESP and services are always originated and terminated on IB-Backbone Edge Bridges (IB-BEBs). IB-BEBs are constituted of I and B components, this is illustrated in Figure 1. A B-component refers to the structure and mechanisms that support the relaying of frames identified by Backbone VLANs in a Provider Backbone Bridge. An I-component refers to the structure and mechanisms that support the relaying of frames identified by service instances (I-SIDs) in a Provider Backbone Bridge. PBB and PBB-TE relay frames with added I-Component TAGs in the I-component and VLAN TAGs in the B-component. PBB and PBB-TE forward frames based on VLAN ID in the VLAN TAG (in the PBB case a B-VID) until the destination MAC address is supported locally by a B-component on this bridge indicating the destination has been reached. At that point, the B-VLAN tag is removed and processing or forwarding on the next TAG begins (in the PBB case an I-Component TAG) until the I-component identified by the I-SID is reached. At the I-component, the I-Component TAG is removed and the next Ethernet type identifies the TAG, etc.

An Ethernet service supported by a PBB-TE TESI is always attached to a Customer Network Port (CNP) of the I-component. A Service Instance Identifier (I-SID) is assigned for the service. I-SIDs are only looked at by source and destination (edge) bridges, so I-SIDs are transparent to path operations and MAY be signaled. The I- and B-components have internal ports that are connected via an internal LAN. These internal ports are the Provider Instance Ports (PIPs) and

Customer Backbone Ports (CBPs). PIPs and CBPs are not visible outside the IB-BEB. ESPs are always originated and terminated on CBP ports and use the MAC address of that port. The I-component encapsulates the service frames arriving from the CNP by adding an I-SID and a complete Ethernet MAC header with an ESP-MAC DA and ESP-MAC SA. The B-component adds the ESP-VID.

This document defines extensions to GMPLS to establish ESPs and TESIs. As can be seen from the above, this requires configuration of both the I- and B-components of the IB-BEBs connected by the ESPs.

In the GMPLS control plane, TE Router IDs are used to identify the IB-BEBs and Backbone Core Bridges (BCBs), and TE Links describe links connected to PNPs and CNPs. TE Links are not associated with CBPs or PIPs.

Note that since multiple internal CBPs may exist, an IB-BEB receiving a PATH message MUST be able to determine the appropriate CBP that is the termination point of the Eth-LSP. To this end, IB-BEBs SHOULD advertise the CNP TE Links in the GMPLS control plane and RSVP-TE signaling SHOULD use the CNP TE Links to identify the termination point of Eth-LSPs. An IB-BEB receiving a PATH message specifying one of its CNPs can locally determine which CBPs have internal connectivity to the I-component supporting the given CNP. In the case that there is more than one suitable CBP, and no I-SID information is provided in the PATH message or previously in the associated Call setup, then the IB-BEB can decide freely which CBP to assign to the requested connection. On the other hand, if there is information on the service (I-SID) that the given ESP will support, then the IB-BEB MUST first determine which PIP and associated CBP is configured with the I-SID and MUST assign that CBP to the ESP.

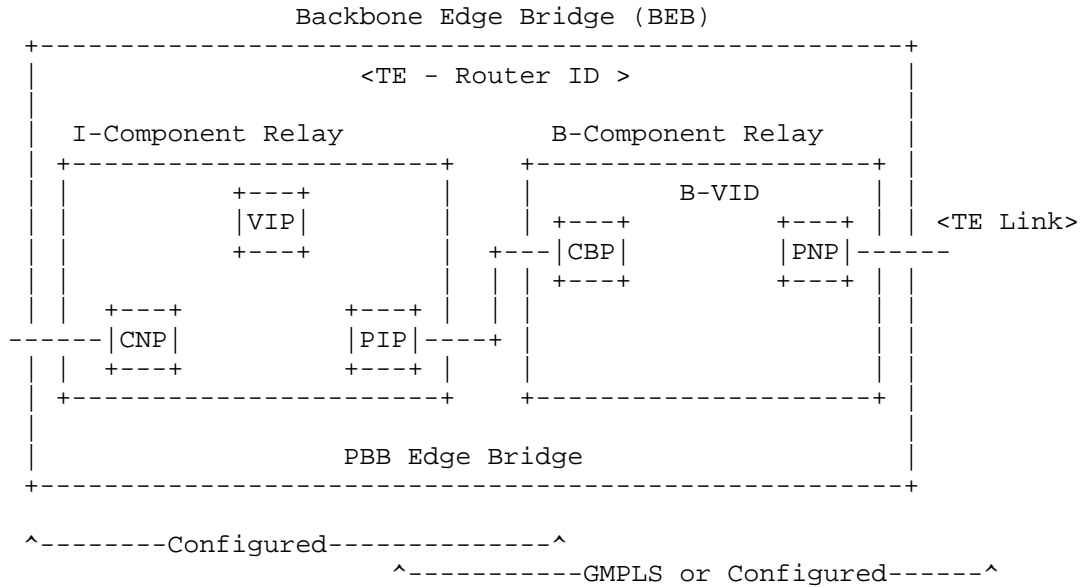


Figure 1: IB-BEBs and GMPLS Identifiers

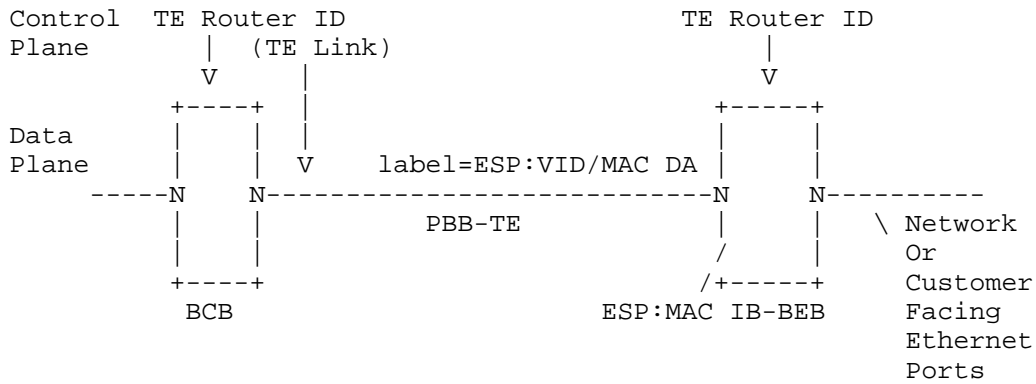


Figure 2: Ethernet/GMPLS Addressing and Label Space

PBB-TE defines the tuple of <ESP-MAC DA, ESP-MAC SA, ESP-VID> as a unique connection identifier in the data plane, but the forwarding operation only uses the ESP-MAC DA and the ESP-VID in each direction. The ESP-VID typically comes from a small number of VIDs dedicated to PBB-TE. ESP-VIDs can be reused across ESPs. There is no requirement that ESP-VIDs for two ESPs that form a P2P TESI be the same.

When configuring an ESP with GMPLS, the ESP-MAC DA and ESP-VID are carried in a generalized label object and are assigned hop by hop, but are invariant within a domain. This invariance is similar to GMPLS operation in transparent optical networks. As is typical with other technologies controlled by GMPLS, the data plane receiver MUST accept, and usually assigns, labels from its available label pool. This, together with the label invariance requirement mentioned above, result in each PBB-TE Ethernet Label being a domain-wide unique label, with a unique ESP-VID + ESP-MAC DA, for each direction.

The following illustrates PBB-TE Ethernet Labels and ESPs for a P2P TESI.

GMPLS Upstream Label	<ESP:MAC1(DA), VID1> (60 bits)
GMPLS Downstream Label	<ESP:MAC2(DA), VID2> (60 bits)
Upstream PBB-TE ESP 3-tuple	<ESP:MAC1, MAC2, VID1> (108 bits)
Downstream PBB-TE ESP 3-tuple	<ESP:MAC2, MAC1, VID2> (108 bits)

Table 1: Labels and ESPs

3.1. Shared Forwarding

One capability of a connectionless Ethernet data plane is to reuse destination forwarding entries for packets from any source within a VLAN to a destination. When setting up P2P PBB-TE connections for multiple sources sharing a common destination, this capability MAY be preserved provided certain requirements are met. We refer to this capability as "shared forwarding". Shared forwarding is invoked based on policy when conditions are met. It is a local decision by label allocation at each end plus the path constraints. Shared forwarding has no impact on the actual paths that are set up, but it allows the reduction of forwarding entries. Shared forwarding paths are identical in function to independently routed paths that share a path from an intersecting bridge or link except they share a single forwarding entry.

The forwarding memory savings from shared forwarding can be quite dramatic in some topologies where a high degree of meshing is required; however, it is typically easier to achieve when the connectivity is known in advance. Normally, the originating GMPLS switch will not have knowledge of the set of shared forwarding paths rooted on the source or destination switch.

Use of a Path Computation Element [RFC4655] or other planning style of tool with more complete knowledge of the network configuration is a way to impose pre-selection of shared forwarding with multiple paths using a single forwarding entry and optimizing for both

directions. In this scenario, the originating bridge uses the LABEL_SET and UPSTREAM_LABEL objects to indicate the selection of the shared forwarding labels at both ends.

3.2. P2P Connections Procedures for Shared Forwarding

The ESP-VID/ESP-MAC DA can be considered to be a shared forwarding identifier or label consisting of some number of P2P connections distinctly identified by the <ESP-MAC DA, ESP-MAC SA, ESP-VID> tuple. This is analogous to an LDP label merge, but in the shared forwarding case, the ESP header contains sufficient information to identify the flow to which a packet belongs. Resources can continue to be allocated per LSP with shared forwarding.

VLAN-tagged Ethernet packets include priority marking. Priority bits MAY be used to indicate Class of Service (COS) and drop priority. Thus, traffic from multiple COSs could be multiplexed on the same Eth-LSP (i.e., similar to E-LSPs) and queuing and drop decisions are made based on the p-bits. This means that the queue selection can be done based on a per-flow basis (i.e., Eth-LSP + priority) and is decoupled from the actual steering of the packet at any given bridge.

A bridge terminating an Eth-LSP will frequently have more than one suitable candidate for sharing a forwarding entry (common ESP-VID/ESP-MAC DA, unique ESP-MAC SA). It is a local decision of how this is performed but a good choice is a path that reduces the requirement for new forwarding entries by reusing common existing paths.

The concept of bandwidth management still applies equally well with shared forwarding.

4. Specific Procedures

4.1. P2P Ethernet LSPs

PBB-TE is designed to be bidirectional and symmetrically routed just like Ethernet. That is, complete and proper functionality of Ethernet protocols is only guaranteed for bidirectional Eth-LSPs. In this section, we discuss the establishment of bidirectional Eth-LSPs.

Note, however, that it is also possible to use RSVP-TE to configure unidirectional ESPs, if the UPSTREAM_LABEL is not included in the PATH message.

To initiate a bidirectional Eth-LSP, the initiator of the PATH message MUST use the procedures outlined in [RFC3473] with the following specifics:

- 1) it MUST set the LSP encoding type to Ethernet (2) [RFC3471].
- 2) it MUST set the LSP switching type to "802_1 PBB-TE", value 40.
- 3) it SHOULD set the Generalized Payload Identifier (G-PID) to Ethernet (33) [RFC3471].
- 4) it MUST set the UPSTREAM_LABEL to the ESP-VID1/ESP-MAC1 tuple where the ESP-VID1 is administered locally for the local MAC address: MAC1.
- 5) it SHOULD set the LABEL_SET or SUGGESTED_LABEL if it chooses to influence the choice of ESP-VID/ESP-MAC DA.
- 6) it MAY carry an I-SID via Call/Connection ID [RFC4974].

Intermediate and egress bridge processing is not modified by this document, i.e., is per [RFC3473]. However, as previously stated, intermediate bridges supporting the 802_1 PBB-TE switching type MUST NOT modify LABEL values.

The ESP-VID1/ESP-MAC1 tuple contained in the UPSTREAM_LABEL is used to create a static forwarding entry in the Filtering Database of bridges at each hop for the upstream direction. This behavior is inferred from the switching type, which is 802_1 PBB-TE. The port derived from the RSVP_HOP object and the ESP-VID1 and ESP-MAC1 included in the PBB-TE Ethernet Label constitute the static entry.

At the destination, an ESP-VID (ESP-VID2) is allocated for the local MAC address: MAC2, the ESP-VID2/ESP-MAC2 tuple is passed in the LABEL object in the RESV message. As with the PATH message, intermediate bridge processing is per [RFC3473], and the LABEL object MUST be passed on unchanged, upstream. The ESP-VID2/ESP-MAC2 tuple contained in the LABEL object is installed in the forwarding table as a static forwarding entry at each hop. This creates a bidirectional Eth-LSP as the PATH and RESV messages follow the same path.

4.1.1. P2P Path Maintenance

Make-before-break procedures can be employed to modify the characteristics of a P2P Eth-LSP. As described in [RFC3209], the LSP ID in the sender template is updated as the new path is signaled. The procedures (including those for shared forwarding) are identical to those employed in establishing a new LSP, with the extended tunnel

ID in the signaling exchange ensuring that double booking of an associated resource does not occur.

Where individual paths in a protection group are modified, signaling procedures MAY be combined with Protection Switching (PS) coordination to administratively force PS operations such that modification is only ever performed on the protection path. PS is a native capability of PBB-TE [IEEE802.1Qay] that can operate when two paths are set up between two common endpoints.

4.2. P2MP Ethernet-LSPs

PBB-TE supports P2MP VID/Multicast MAC (MMAC) forwarding. In this case, the PBB-TE Ethernet Label consists of a VID and a Group MAC address. The procedures outlined in [RFC3473] and [RFC4875] could be adapted to signal P2MP LSPs for the source (point) to destination (multipoint) direction. Each one of the branches of the P2MP Eth-LSP would be associated with a reverse-path symmetric and congruent P2P Eth-LSP.

Complete procedures for signaling bidirectional P2MP E-LSPs are out of scope for this document.

4.3. PBB-TE Ethernet Label

The PBB-TE Ethernet Label is a new generalized label with the following format:

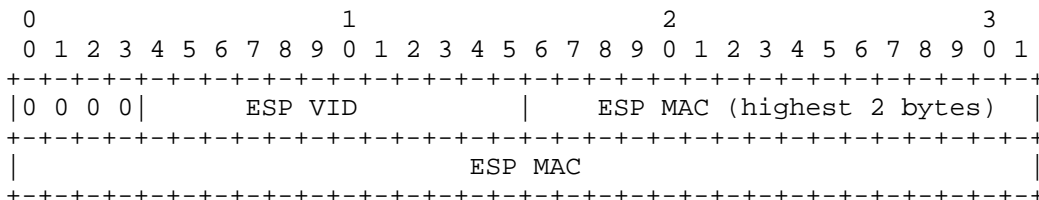


Figure 3: PBB-TE Ethernet Label

This format MUST be used for both P2P and P2MP Eth-LSPs. For P2P Eth-LSPs, the fields specify a VID and a unicast MAC address; whereas, for P2MP Eth-LSPs, a VID and a group MAC address is carried in the label. The PBB-TE Ethernet Label is a domain-wide unique label and MUST be passed unchanged at each hop. This has similarity to the way in which a wavelength label is handled at an intermediate bridge that cannot perform wavelength conversion, and is described in [RFC3473].

4.4. Protection Paths

When protection is used for path recovery, it is required to associate the working and protection paths into a protection group. This is achieved as defined in [RFC4872] and [RFC4873] using the ASSOCIATION and PROTECTION objects.

4.5. Service Instance Identification

The I-SID is used to uniquely identify services within the network. Unambiguous identification is achieved by ensuring global uniqueness of the I-SIDs within the network or at least between any pair of edge bridges. On IB-BEBs, the Backbone Service Instance Table is used to configure the mapping between I-SIDs and ESPs. This configuration can be either manual or semi-automated by signaling described here.

RSVP-TE Signaling MAY be used to automate I-SID to ESP mapping. By relying on signaling, it is ensured that the same I-SID is assigned to the service and mapped to the same ESP. Note, by signaling the I-SID associated to the ESP, one can ensure that IB-BEBs select the appropriate CBP port.

CALL signaling [RFC4974] MAY be used to create an association between the Eth-LSP endpoints prior to establishment of the LSP. The CALL_ATTRIBUTES object can be used during CALL signaling, as described in [RFC4974], to indicate properties of the CALL. The Service ID TLV, defined below, can be carried in the CALL_ATTRIBUTES object to indicate the I-SID to ESP mapping for the Eth-LSP that will be set up in association with the CALL.

Alternatively, the GMPLS RSVP-TE PATH message can carry the I-SID association using the Service ID TLV in the LSP_ATTRIBUTES object [RFC5420] at the time of Eth-LSP signaling. Using this mechanism, it is possible to create the I-SID association, either when the path is set up or at a later time using a PATH refresh.

A new Service ID TLV is defined for the CALL_ATTRIBUTES and LSP_ATTRIBUTES objects. The type value is 3 when carried in the CALL_ATTRIBUTES object and the type value is 2 when carried in the LSP_ATTRIBUTES object. The format is depicted below.

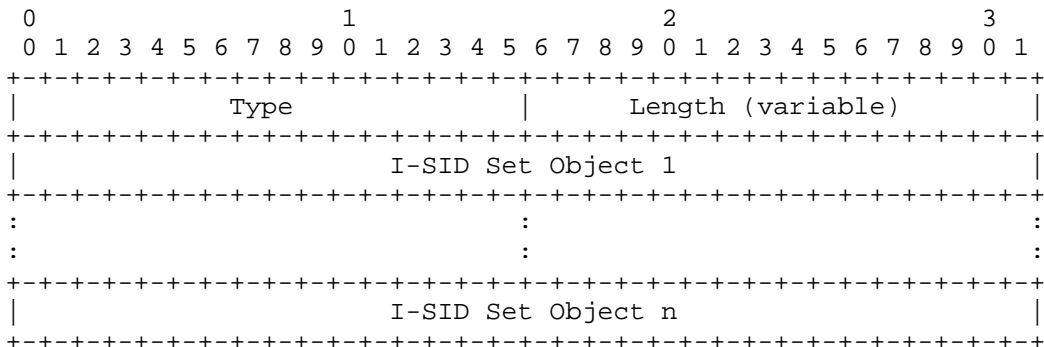


Figure 4: Service ID TLV

- I-SID Set Object: is used to define a list or range of I-SIDs. Multiple I-SID Set Objects can be present. At least one I-SID Set Object MUST be present. In most of the cases, a single I-SID Set Object with a single I-SID value is used. The I-SID Set Object is used to define a list or range of I-SIDs. The format of the I-SID Set Object is based on the LABEL_SET Object:

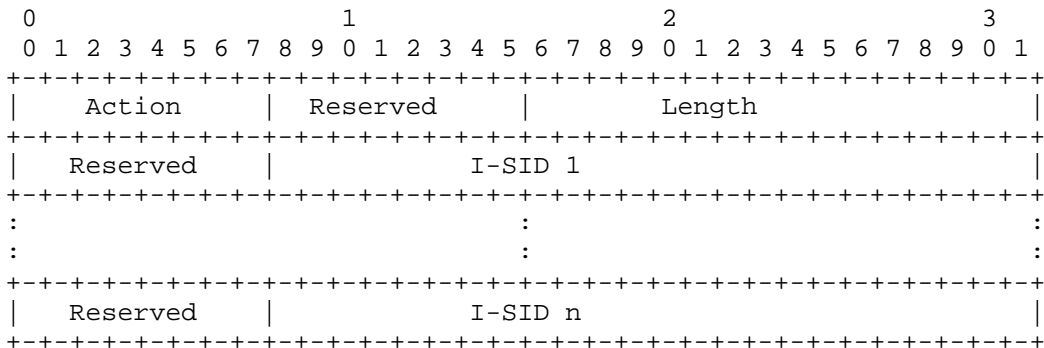


Figure 5: I-SID Set Object

- Action: 8 bits

The following actions are defined: list (0), range (1). When a range is defined, there are only two I-SIDs that follow the beginning I-SID and the end of the range I-SID. When list is defined, a number of I-SIDs may be defined.

- Length: 16 bits

This indicates the length of the I-SID Set object.

- I-SID: 24 bits

The I-SID value identifies a particular backbone service instance.

5. Error Conditions

The following errors identify Eth-LSP-specific problems.

In PBB-TE, a set of ESP-VIDs allocated to PBB-TE must be configured. Therefore, it is possible in some situations that the configuration of a bridge is not the same as other bridges. If the ESP-VIDs of various bridges have some ESP-VIDs in common, it is possible some paths may be set up before encountering issues. This is a management issue since all bridges should have the same ESP-VID range. Configuration should be consistent.

5.1. ESP-VID-Related Errors

The network operator administratively selects a set of VLAN Identifiers that can be used to set up ESPs. Consequently, any VID outside the allocated range is invalid, and an error MUST be generated where the mismatch is discovered. The Error indication is carried in the PathErr message from any intermediate bridge that does not support the signaled source VID or optionally the destination VID. The Error MAY be indicated in the ResvErr if the allocation error happens on the RESV message. In this case, a bridge that does not support the signaled destination VID MUST signal the error.

5.1.1. Invalid ESP-VID Value in the PBB-TE Ethernet Label

If a bridge is not configured to use the ESP-VID value, carried in the Label object, for PBB-TE ESPs, it MUST immediately generate an error: Routing problem (24) / Unacceptable label value (6). Handling of this error is according to [RFC3209].

Note that an originating bridge can reuse an ESP-VID with a different source or destination B-MAC address. By allocating a number of B-MACs and a number of ESP-VIDs, a large number of PBB-TE connections may be supported.

Note, this error may be originated by any bridge along the path.

5.1.2. Allocated ESP-VID Range is Exhausted

The destination bridge, after receiving the PATH message, has to assign a VID, which, together with its MAC address, will constitute the PBB-TE Ethernet Label. An existing VID may be reused when shared forwarding is used or when there are no path conflicts; otherwise, the bridge has to allocate a VID.

Depending on the size of the allocated VLAN range and the number of Eth-LSPs terminated on a particular bridge, it is possible that the available VIDs are exhausted; hence, no PBB-TE Ethernet Label can be allocated. In this case, the destination bridge SHOULD generate a PathErr message with error code: Routing problem (24) and error value: MPLS Label allocation failure (9).

5.2. Invalid MAC Address

IEEE defines a set of reserved MAC addresses from 01-80-C2-00-00-00 to 01-80-C2-00-00-0F as explained in [IEEE802.1Q] that have special meaning, processing, and follow specific forwarding rules. These addresses cannot be used for PBB-TE ESPs. In the case the PBB-TE Ethernet Label refers to such a MAC address, a bridge encountering the mismatch MUST immediately generate an error: Routing problem (24) / Unacceptable label value (6). Handling of this error is according to [RFC3209].

6. Security Considerations

This document does not introduce new security issues; the considerations in [RFC4872] and [RFC4873] apply.

A GMPLS-controlled Ethernet PBB-TE system assumes that users and devices attached to User-to-Network Interfaces (UNIs) may behave maliciously, negligently, or incorrectly. Intra-provider control traffic is trusted not to be malicious. In general, these requirements are no different from the security requirements for operating any GMPLS network. Access to the trusted network will only occur through the protocols defined for the UNI or Network-to-Network Interface (NNI) or through protected management interfaces.

When in-band GMPLS signaling is used for the control plane, the security of the control plane and the data plane may affect each other. When out-of-band GMPLS signaling is used for the control plane, the data-plane security is decoupled from the control plane; therefore, the security of the data plane has less impact on overall security.

Where GMPLS is applied to the control of VLAN only, the commonly known techniques for mitigation of Ethernet denial-of-service (DoS) attacks may be required on UNI ports. PBB-TE has been designed to interwork with legacy VLANs and the VLANs provide isolation from Ethernet legacy control planes.

Where control-plane communications are point-to-point over links that employ 802.1AE Media Access Control Security [MACSEC], it may reasonably be determined that no further security measures are used. In other cases, it is appropriate to use control-plane security where it is deemed necessary to secure the signaling messages. GMPLS signaling security measures are described in [RFC3471] and [RFC3473], and they inherit security techniques applicable to RSVP-TE, as described in [RFC3209] and [RFC2205]. For a fuller overview of GMPLS security techniques, see [RFC5920].

7. IANA Considerations

A new Switching Type, "802_1 PBB-TE" (40), has been assigned in the Switching Types registry of the GMPLS Signaling Parameters registry.

The Service ID TLV has been assigned in the Attributes TLV Space in the RSVP-TE Parameters registry. It is carried in the LSP_ATTRIBUTES object (class = 197, C-Type = 1) [RFC5420]. This new type has been registered as follows:

```
Type: 2
Name: Service ID TLV
Allowed on LSP_ATTRIBUTES: Yes
Allowed on LSP_REQUIRED_ATTRIBUTES: No
```

The Service ID TLV has been assigned value 3 in the Call Attributes TLV registry in the RSVP Parameters registry. It is carried in the CALL_ATTRIBUTES object (class = 202, C-Type = 1) defined by [RFC6001].

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.
- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4872] Lang, J., Ed., Rekhter, Y., Ed., and D. Papadimitriou, Ed., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, May 2007.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.
- [RFC4974] Papadimitriou, D. and A. Farrel, "Generalized MPLS (GMPLS) RSVP-TE Signaling Extensions in Support of Calls", RFC 4974, August 2007.
- [RFC5420] Farrel, A., Ed., 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.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiomoto, K., Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 6001, October 2010.

8.2. Informative References

- [IEEE802.1ah]
"IEEE Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks - Amendment 6: Provider Backbone Bridges", (2008)
- [IEEE802.1Q]
"IEEE Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2005, May 19, 2006.
- [IEEE802.1Qay]
"IEEE Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks - Amendment : Provider Backbone Bridges Traffic Engineering", 2009.
- [MACSEC] "IEEE Standard for Local and metropolitan area networks Media Access Control (MAC) Security", IEEE 802.1AE-2006, August 18, 2006.
- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [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., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

9. Acknowledgments

The authors would like to thank Dinesh Mohan, Nigel Bragg, Stephen Shew, Dave Martin and Sandra Ballarte for their contributions to this document. The authors thank Deborah Brungard and Adrian Farrel for their review and suggestions to this document.

Authors' Addresses

Don Fedyk
Alcatel-Lucent
Groton, MA 01450
Phone: +1-978-467-5645
EMail: donald.fedyk@alcatel-lucent.com

Himanshu Shah
Ciena
1741 Technology Dr, #400
San Jose, CA 95110
Phone: 508-435-0448
EMail: hshah@ciena.com

Nabil Bitar
Verizon
40 Sylvan Rd.
Waltham, MA 02451
EMail: nabil.n.bitar@verizon.com

Attila Takacs
Ericsson
1. Laborc u.
Budapest, HUNGARY 1037
EMail: attila.takacs@ericsson.com