

Network Working Group
Internet Draft
Category: Standards Track

Don Fedyk, David Allan, Nortel
Himanshu Shah, Ciena
Nabil Bitar, Verizon
Attila Takacs, Diego Caviglia, Ericsson
Alan McGuire, BT
Nurit Sprecher, Nokia Siemens Networks
Lou Berger, LabN
November 19, 2007

GMPLS control of Ethernet PBB-TE
draft-fedyk-gmpls-ethernet-pbb-te-02.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire in May 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This memo is complementary to [\[ARCH\]](#) and describes how a GMPLS control plane may be applied to the Provider Backbone Bridges Traffic Engineering (PBB-TE) [IEEE 802.1Qay] amendment to 802.1Q and how GMPLS can be used to configure VLAN-aware Ethernet switches in order to establish Ethernet point to point (P2P) and P2MP MAC switched

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

paths and P2P/P2MP VID based trees. This document supports, but does not modify, the standard IEEE data.

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](#)].

Document History

This document has under gone name changes to follow the standardization of Provider Backbone Bridges and the Traffic engineering capability.

[draft-fedyk-gmpls-ethernet-ivl-00.txt](#).

This was the original draft.

[draft-fedyk-gmpls-ethernet-pbt-00.txt](#)
[draft-fedyk-gmpls-ethernet-pbt-01.txt](#)

This draft was renamed to reflect the Provider Backbone Transport (PBT) nomenclature. Several co-authors joined the draft.

[draft-fedyk-gmpls-ethernet-pbb-te-00.txt](#)

The standardization of PBT is called Provider Backbone Bridges Traffic Engineering (PBB-TE). The draft was aligned the PBB-TE Technology.

[draft-fedyk-gmpls-ethernet-pbb-te-01.txt](#)

This is the second revision of the PBB-TE draft with editing to clarify the document and the addition of co-authors.

[draft-fedyk-gmpls-ethernet-pbb-te-02.txt](#)

This is a third revision with the general aspects of Ethernet being move to the architecture and framework [[ARCH](#)] and the specifics for PBB-TE becoming more clear.

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

Table of Contents

1. Introduction.....	4
2. Terminology.....	4
2.1 PBB-TE Terminology.....	5
3. GMPLS creation and maintenance of PBB-TE Service Instances....	5
3.1 Ethernet Service.....	6
3.2 Addresses, Interfaces, and Labels.....	7
4. Specific Procedures.....	9
4.1 P2P connections.....	9
4.1.1 Shared Forwarding.....	10
4.1.2 P2P connections with shared forwarding.....	11
4.1.3 Dynamic P2P symmetry with shared forwarding.....	12
4.1.4 Planned P2P symmetry.....	12
4.1.5 P2P Path Maintenance.....	12
4.2 P2MP Signaling.....	13
4.3 P2MP VID/ESP-MAC DA Connections.....	13
4.3.1 Setup procedures.....	13
4.3.2 Maintenance Procedures.....	13
4.4 Ethernet Label.....	14
4.5 OAM MEP ID and MA ID synchronization.....	15
4.6 Protection Paths.....	15
5. Error conditions.....	16
5.1 Invalid ESP-VID value for PBB-TE MSTI range.....	16
5.2 Invalid MAC Address.....	16
5.3 Invalid ERO for UPSTREAM_LABEL Object.....	16
5.4 Invalid ERO for LABEL_SET Object	16
5.5 Switch is not ESP P2MP capable.....	16
5.6 Invalid ESP-VID in UPSTREAM_LABEL object	16
6. Deployment Scenarios.....	16
7. Security Considerations	16
8. IANA Considerations.....	17
9. References.....	17
9.1 Normative References.....	17

9.2	Informative References.....	17
10.	Author's Address.....	18
11.	Intellectual Property Statement.....	19
12.	Disclaimer of Validity.....	20
13.	Copyright Statement.....	20
14.	Acknowledgments.....	20
Appendix A	21
	Rational and mechanism for PBB_TE Ethernet Forwarding.....	21
A 1.	Overview of configuration of VID/DMAC tuples.....	23
A 2.	Overview of configuration of VID port membership.....	26
A 3.	OAM Aspects	26
A 4.	QOS Aspects	27
A 5.	Resiliency Aspects	27
A 5.1.	E2E Path protection.....	27

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

[1.](#) Introduction

IEEE 802.1 is specifying Traffic Engineered Ethernet paths in the Provider Backbone Bridged network (PBB-TE) [IEEE 802.1Qay] based on managed objects that can be separated from the Spanning Tree Control Plane and statically configured or managed by a another control plane. These paths have minor changes to Ethernet data plane specified in the IEEE. IEEE 802 termed these paths "PBB-TE service instances".

The purpose of this document is to specify extensions for a GMPLS based control plane to manage PBB-TE service instances. This draft is aligned with GMPLS Ethernet Label Switching Architecture and Framework [[ARCH](#)].

It should be noted that due to the changes in the separation of the Spanning Tree Control plane and the PBB-TE forwarding, the behavior of PBB-TE for the specified VLAN range is a new behavior. (It does not default to conventional Ethernet forwarding with learning at any time). [Appendix A](#) summarized the rational for this data plane technology until the IEEE specification is more mature.

[2.](#) Terminology

In addition to well understood GMPLS terms, this memo uses

terminology from IEEE 802.1 and introduces a few new terms:

B-MAC	Backbone MAC
B-VID	Backbone VLAN ID
B-VLAN	Backbone VLAN
CBP	Customer Backbone Port
CCM	Continuity Check Message
COS	Class of Service
CLI	Command Line Interface
CIP	Customer Instance Port
C-MAC	Customer MAC
C-VID	Customer VLAN ID
C-VLAN	Customer VLAN
DMAC	Destination MAC Address
ESP	Ethernet Switched Path
Eth-LSP	Ethernet Label switched Path
I-SID	Ethernet Service Instance Identifier
LBM	Loopback Message
LBR	Loopback Reply
LLDP	Link Layer Discovery Protocol
LMM	Loss Measurement Message
LMR	Loss Measurement Reply
MAC	Media Access Control
MMAC	Multicast MAC
MSTI	Multiple Spanning Tree Instance
MP2MP	Multipoint to multipoint
PBB	Provider Backbone Bridges
PBB-TE	Provider Backbone Bridges Traffic Engineering

Fedyk et al.

Expires May 2008

[Page 4]

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

PIP	Provider Instance Port
PNP	Provider Network Port
P2P	Point to Point
P2MP	Point to Multipoint
QOS	Quality of Service
ESP-MAC SA	Source MAC Address
S-VID	Service VLAN ID
SVL	Shared VLAN Learning
VID	VLAN ID
VLAN	Virtual LAN

[2.1](#) PBB-TE Terminology

The PBB-TE specification has defined some additional terminology to

clarify the PBB-TE functions. We repeat these here in expanded context to translate from IEEE to GMPLS terminology.

- Ethernet Switched Path (ESP): A provisioned traffic engineered unidirectional connectivity path between two or more Customer Backbone Ports (CBPs) which extends over a Provider Backbone Bridge Network (PBBN). The path is identified by the 3-tuple <ESP-MAC DA, ESP-MAC SA, ESP-VID> where the ESP-VID value is allocated to the PBB-TE Multiple Spanning Tree Instance (MSTI) (A set of VIDs for PBB-TE is allocated as a set of MSTIs). An ESP is analogous to an GMPLS LSP.
- PBB-TE Region: A set of PBB switches and PB switches by a set of Service-VLANs allocated to provisioned Ethernet Switched Paths (ESPs).
- PBB-TE service instance: A Point-to-Point or a Point-to-Multipoint PBB-TE service instance.
- PBB-TE Trunk: A Point-to-Point PBB-TE service instance.
- Point-to-Point PBB-TE service instance: An instance of the MAC service provided by two unidirectional co-routed ESPs forming a bidirectional service. A GMPLS bidirectional path is analogous to a P2P PBB-TE Service instance.
- Point-to-Multipoint PBB-TE service instance: An instance of the MAC service provided by a set of ESPs which comprises one multipoint ESP plus n unidirectional point-to-point ESPs, routed along the leaves of the multicast ESP. A P2MP GMPLS bidirectional tree is analogous to a P2MP PBB-TE service instance.

3. GMPLS creation and maintenance of PBB-TE Service Instances

PBB-TE is an Ethernet connection oriented technology, being specified in the IEEE, which can be controlled by configuration of static filtering entries [see [Appendix A](#)] for some details on the

rational for the data plane. PBB-TE ESPs are created switch by switch by simple configuration of Ethernet logical ports and assignment of PBB-TE labels or by a control plane. This document describes GMPLS as a valid control plane for Eth-LSPs that are based on PBB-TE ESPs. A Point-to-Point PBB-TE service instance is a form of Ethernet LSP (Eth-LSP) which is more broadly defined in [[ARCH](#)]. This memo describes GMPLS as a mechanism to automate set-up

teardown, protection and recovery of PBB-TE ESPs and specifies the specific TLVs for control of PBB-TE service instances.

When configuring a PBB-TE 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 label being a domain wide unique label, with a unique ESP-VID + ESP-MAC DA, for each direction.

The following illustrates the identifiers for Labels and ESPs.

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

The MAC is domain wide unique in the network. 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 (DMAC) and the ESP-VID in each direction. Note that the MAC addresses for PBB-TE are part of the Backbone Component Relay (B-Component) and are associated with Provider addresses corresponding to the Backbone Customer ports as described in [section 3.2](#). The ESP-VID (VID) typically comes from a small number of VIDs dedicated to PBB-TE MSTI. The ESP-VID (VID) can be reused across ESPs. There is no requirement the ESP-VID for two ESPs that for a PBB-TE Service instance be the same.

Several attributes may be associated with an Eth-LSP. These are reviewed in Section 3 of [\[ARCH\]](#). Several other aspects of GMPLS covered by [\[ARCH\]](#) also apply equally to PBB-TE. This includes the GMPLS routing and addressing model, link management, path computation and selection, and multiple domains.

[3.1](#) Ethernet Service

Ethernet Switched Paths that are setup either by configuration or signaling can be used to provide an Ethernet service to customers of the Ethernet network. The Metro Ethernet Forum has defined some

services in [MEF.6] (e.g., Ethernet Private Line), and these are also aligned with ITU-T G.8011-x Recommendations. Of particular interest are the bandwidth profile parameters in [MEF.10] and whose associated bandwidth profile algorithm are based on [RFC4115] [RFC3270]. Consideration should be given to supporting these in any signaling extensions for Ethernet LSPs. This will be addressed in a future version of this specification.

[3.2](#) Addresses, Interfaces, and Labels

This specification uses an addressing scheme and a label space for the ingress/egress connection; the hierarchical TE Router ID/Interface ID and the Ethernet ESP-VID/ESP-MAC DA tuple or ESP-VID/Multicast MAC as a label space. This draft is intended to be consistent with GMPLS addressing and Routing [ARCH].

PBB-TE is defined for a PBB IB-Bridge. This is illustrated in Figure 1. The Ethernet service is attached to a Customer Instance Port (CIP) of the Backbone Service Instance (I-component) Relay. The CIP is interfaced to a Virtual instance port (VIP) which is identified with a configured service instance (I-SID) and attached to a Provider Instance Port (PIP). The PIP is configured to be attached to a customer Backbone port (CBP). (A point to point service instance is illustrated. A point to multipoint service could allow more than one CBP to be attached to a single PIP.) The CBP has a BMAC that defines the MAC for the PBB-TE Service Instance. The B-Component relay adds the ESP Header the ESP-MAC DA, ESP-MAC SA and the ESP-VID. GMPLS is being defined here to connect CPB MACs to signal the PBB-TE service Instance before the association of ESP-MAC DA and ESP-MAC SA is defined.

The diagram also shows the addition of a TE Router ID to the PBB switch and the TE Link identifier to enable GMPLS. TE Links are not associated with CPBs. TE Links are associated with PNPs. TE links are associated with node identifiers of backbone edge bridges (BEB) and backbone core bridges (BCB). CBPs are also associated with these node ids. For GMPLS the node IDs are expressed as IP addresses as TE-Router IDs. [ADDRESS]

Internet Draft

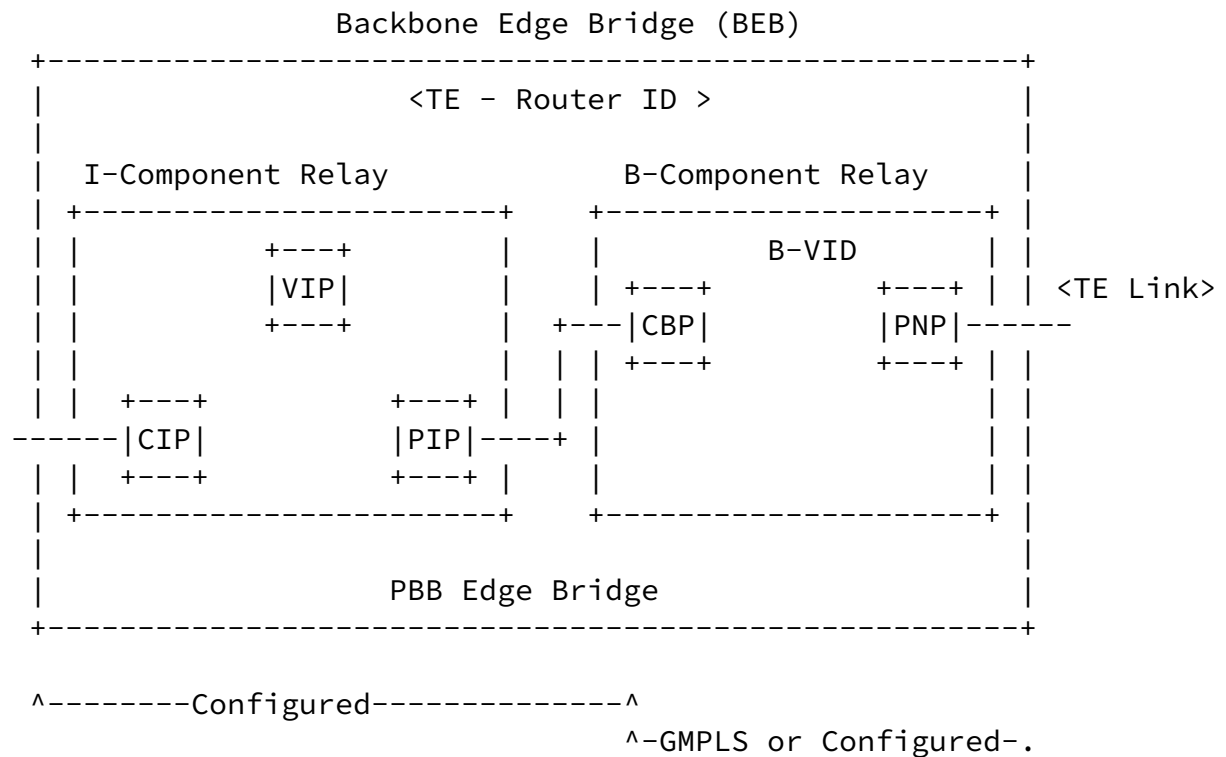
[draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

Figure 2 Ethernet/GMPLS Addressing & Label Space

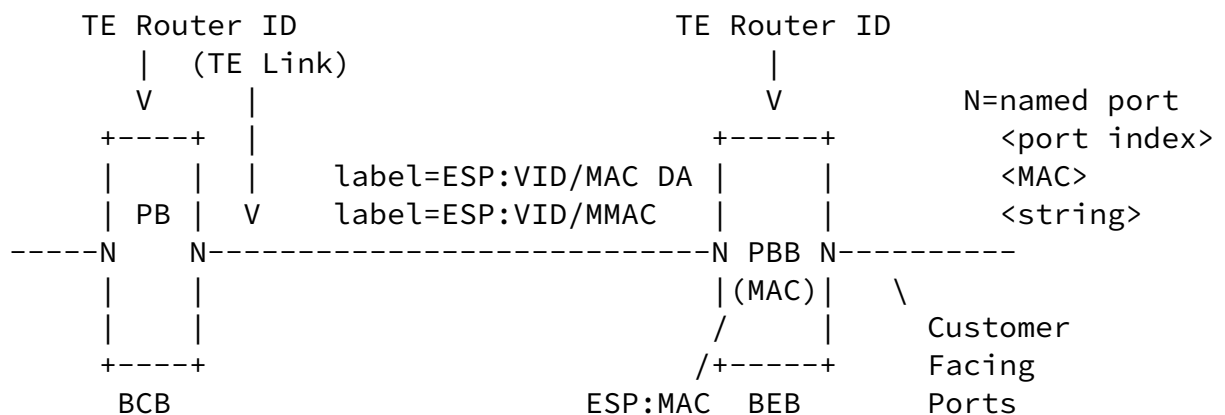


Figure 3 Ethernet/GMPLS Addressing & Label Space

For a GMPLS based system, the TE Router ID/logical port is the logical signaling identifier for the control plane via which Ethernet layer label bindings are solicited. In order to create a P2P path an association must be made between the ingress and egress node. The actual label distributed via signaling and instantiated in the switch forwarding tables identifies the upstream and downstream egress ESP-VID/ESP-MAC DA of the PBB-TE ESP (see Figure 4).

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

GMPLS uses identifiers in the form of 32 bit numbers which are in the IP address notation which may or may not be IP addresses. The provider MAC port addresses are exchanged by the LLDP [IEEE 802.1AB] and by LMP [RFC4204] if supported. However these identifiers are merely for link control and legacy Ethernet support and have local link scope. Actual label assignment is performed by the ingress and egress nodes using CPB MAC addresses.

A particular PNP would have:

- A VID/MAC
- An IP address, which is typically the TE router ID, plus a 32 bit interface Identifier also call an unnumbered link.
- One (or more) Mnemonic String Identifiers

This multiple naming convention leaves the issue of resolving the set given one of the port identifiers. On a particular node, mapping is relatively straightforward. Per [ARCH], standard GMPLS mechanisms are used for signaling resolution. In so doing, the problem of setting up a path is reduced to figuring out what switch supports an egress CBP MAC address and then finding the corresponding egress IP address and performing all signaling and routing with respect to the egress.

There are several options to achieve this:

- Provisioning
- Auto discovery protocols that carry MAC address
- Augmenting Routing TE with MAC Addresses
- Name Servers with identifier/address registration

The specific procedures will be clarified in a subsequent version of this document.

[4. Specific Procedures](#)

[4.1 P2P connections](#)

The PBB-TE Service Instance is defined by the ESP 3-tuples for each of the unidirectional ESPs. From a GMPLS control plane point of view an Ethernet LSP MAY also be identified as any other LSP using the 5-tuple [Ip_Source_Saddr, Ip_Dest_Addr, LSP_Id, Tunnel_ID, Extended_Tunnel_ID]. The ESP-VID and ESP-MAC DA tuple identifies the forwarding multiplex at transit switches and a simple degenerate form of the multiplex is a single P2P connection.

This results in unique labels end to end. The data streams MAY merge, the forwarding entries MAY be shared but the headers are still unique allowing the connection to be de-multiplexed downstream.

On the initiating and terminating nodes, a function administers the ESP-VIDs associated with the ESP-MAC SA and ESP-MAC DA respectively. PBB-TE is designed to be bidirectional and symmetrically routed just like Ethernet. Therefore in PBB-TE, the packet ESP-MAC SA and ESP-

MAC DA pair is same in the forwarding path and the associated reverse path except they are flipped in the reverse direction.

To initiate a bidirectional ESP-VID/ESP-MAC DA P2P or P2MP path, the initiator of the PATH message uses procedures outlined in [[RFC3473](#)] possibly augmented with [[RFC4875](#)], it:

- 1) Sets the LSP encoding type to Ethernet.
- 2) Sets the LSP switching type to 802_1 PBB-TE [IANA to define].
- 3) Sets the GPID to service type [IANA to define].
- 4) Sets the UPSTREAM_LABEL to the ESP-VID/ESP-MAC SA tuple where the ESP-VID is administered from the configured ESP-VID/ESP-MAC DA range.
- 5) Optionally sets the LABEL_SET or SUGGESTED_LABEL if it chooses to influence the choice of ESP-VID/ESP-MAC DA.
- 6) Optionally look at Call / Connection ID for Carrying I-SID.

Intermediate and egress node processing is not modified by this document, i.e., is per [\[RFC3473\]](#) and, in the case of P2MP, as extended in [\[RFC4875\]](#). Note, as previously stated intermediate nodes supporting the 802_1 switching type may not modify LABEL values.

The ESP-VID/ESP-MAC SA 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 [IANA to define]. The port derived from the RSVP_HOP object and the ESP-VID and ESP-MAC DA included in the label constitute the static entry.

At the destination, a ESP-VID is allocated in the local MAC range for the ESP-MAC DA and the ESP-VID/ESP-MAC DA tuple is passed in a LABEL object in the RESV message. As with the Path message, intermediate node processing is per [\[RFC3473\]](#) and [\[RFC4875\]](#), and the LABEL object is passed on unchanged, upstream. The ESP-VID/ESP-MAC DA tuple contained in the LABEL Object is installed in the forwarding table as a static forwarding entry at each hop. This creates a bidirectional path as the PATH and RESV messages follow the same path.

[4.1.1](#) 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. Shared forwarding has no impact on the actual paths setup, but it allows the reduction of forwarding entries. Shared forwarding paths are identical to independently routed paths with the exception that they share the same labels and same path from the merge point.

To achieve shared forwarding, a Path computation engine [\[PATHCOMP\]](#) should ensure the ERO is consistent with an existing path for the shared segments. If a path satisfies the consistency check, the upstream end of the signaling may chose to share an existing ESP-VID/ESP-MAC DA for the upstream traffic with an existing Eth-LSP.

The criteria for shared forwarding is the Eth-LSPs must share the same destination port and the paths of the Eth-LSP share one or more hops consecutively. Once the paths converge they must remain converged. If no existing path has this behavior when a new path is being created, the new path will be created without sharing either by using another ESP-VID or another ESP-MAC DA or both.

In other words, shared forwarding is possible when paths share segments either from the source or the destination. There is no requirement that the paths share reservations or other attributes. For the source, the UPSTREAM_LABEL is chosen to be the same as an existing path that shares the ERO for some number of hops. Similarly for the destination, shared forwarding is possible when an existing path that shares segments with the new paths ERO, viewed from the destination switch. The downstream label in this case is chosen to be the same as the existing path. In this manner shared forwarding is a function that is controlled primarily by policy and in combination with the local label allocation at the end points of the path.

[4.1.2](#) P2P connections with shared forwarding

The ESP-VID/ESP-MAC DA MAY be considered to be a shared forwarding identifier or label for a multiplex consisting of some number of P2P connections distinctly identified by the MAC ESP-VID/ESP-MAC DA/ESP-MAC SA tuple. In some ways this is analogous to an LDP label merge but in the shared forwarding case only the forwarding entry is reused. Resources can continue to be allocated per LSP.

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 (i.e., Eth-LSP + priority) basis and is decoupled from the actual steering of the packet at any given node.

A switch terminating an Eth-LSP will frequently have more than one suitable candidate path and it may choose to share a forwarding entry (common ESP-VID/ESP-MAC DA, unique ESP-MAC SA). It is a local

decision of how this is performed but the best choice is a path that maximizes the shared forwarding.

The concept of bandwidth management still applies equally well with shared forwarding. As an example consider a PBB-TE edge switch that terminates an Ethernet LSP with the following attributes: bandwidth B1, ESP-MAC DA D, ESP-MAC SA S1, ESP-VID V. A request to establish an additional Ethernet LSP with attributes (bandwidth B2, ESP-MAC DA D, ESP-MAC SA S2, ESP-VID V) can be accepted provided there is sufficient link capacity remaining.

[4.1.3](#) Dynamic P2P symmetry with shared forwarding

Similar to how a destination switch MAY select a ESP-VID/ESP-MAC DA from the set of existing shared forwarding multiplexes rooted at the destination node, the originating switch MAY also do so for the reverse path. Once the initial ERO has been computed and the set of existing Ethernet LSPs that include the target ESP-MAC DA have been pruned, the originating switch may select the optimal (by whatever criteria) existing shared forwarding multiplex for the new destination to merge with and offer its own ESP-VID/ESP-MAC DA tuple for itself as a destination.

[4.1.4](#) Planned P2P symmetry

Normally the originating switch will not have knowledge of the set of shared forwarding paths rooted on the destination node.

Use of a Path Computation Server [[PATHCOMP](#)] or other planning style of tool with more complete knowledge of the network configuration may wish to impose pre-selection of shared forwarding multiplexes to use for both directions. In this scenario the originating switch uses the LABEL_SET and UPSTREAM_LABEL objects to indicate complete selection of the shared forwarding multiplexes at both ends. This may also result in the establishment of a new ESP-VID/ESP-MAC DA path as the LABEL_SET object may legitimately refer to a path that does not yet exist.

[4.1.5](#) P2P Path Maintenance

Make before break procedures can be employed to modify the characteristics of a P2P Ethernet 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 the associated resources 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 switching operations such that modifications are only ever performed on the protection path.

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

[4.2](#) P2MP Signaling

Note specifics for P2MP paths are being defined. This section will be updated to align with the PBB-TE specification [IEEE 802.1Qay].

To initiate a P2MP VID/Multicast MAC (MMAC) path the initiator of the PATH message uses procedures outlined in [[RFC3473](#)] and [[RFC4875](#)]. A P2MP tree consists of a VID tree or MMAC tree in the forward direction (from root to leaves) and a set of P2P paths running on identical paths from Tree to root in the reverse direction. The result is a composite path with Multicast VID/ESP-MMAC DA labels with a single ESP-MMAC DA in the forward direction and a symmetric unidirectional ESP-VID/ESP-MAC DA label in the reverse direction:

- 1-4) Same points as P2P paths previously specified.
- 5) Sets the downstream label as the Multicast VID/ESP-MMAC DA.
- 6) VID translation may optionally be permitted on a local basis between two switches by a downstream switch replying with a Multicast VID/ESP-MMAC DA other than the LABEL_SET. The upstream switch then sets a VID translation on the port associated with the label to allow VID translation. This flexibility allows the tree to be constructed with out having to worry about colliding with another tree using the same VID. (Inclusion of this point is TBD by [IEEE 802.1Qay])

[4.3](#) P2MP VID/ESP-MAC DA Connections

[4.3.1](#) Setup procedures

The group ESP-MMAC DA is administered from a central pool of multicast addresses and the VLAN selected from the PBB-TE MSTI range. The P2MP tree is constructed via incremental addition of leaves to the tree in signaling exchange where the root is the originating switch (as per ([RFC4875](#))). The multicast VID/ESP-MAC DA is encoded in the LABEL_SET (as a member of one) object using the Ethernet label encoding.

Where a return path is required the unicast MAC corresponding to the originating interface and a VID selected from the configured VID/ESP-MAC DA range is encoded as an Ethernet label in the UPSTREAM_LABEL

is not desired.

These domain wide labels are allocated to switches that control the assignment of labels. There are two options for Ethernet MAC based domain wide unique labels. One option is to allocate the ESP-MAC DAs from globally unique addresses assigned to the either the switch manufacturer or the owner. The other option is to use ESP-MAC DAs out of the local admin space and ensure these labels are unique within the domain. This local ESP-MAC DA space does not have to be globally unique because the labels are only valid within a single provider domain.

In the case of local label allocation there is less administrative overhead to allocate labels. However when using configuration, a tool would have to perform a consistency check to make sure that labels were unique. When using GMPLS signaling it is assumed a unique pool of labels would be assigned to each switch. The ESP-MAC DA addresses are domain wide unique and so is the combination of ESP

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

<VID, MAC DA>. It is intended that the ESP <VID, MAC DA> be only used by one destination. However, should an error occur and somehow a duplicate label be assigned to one or more destination switches GMPLS signaling procedures would allow the first assignment of the label and prevent any duplicate label from colliding. If a collision occurs an alarm would be generated. In fact some of these procedures have been defined in GMPLS control of photonic networks where a lambda may exist as a form of domain wide label.

[4.5](#) OAM MEP ID and MA ID synchronization

This section is aligned with [IEEE 802.1Qay]. At present it Ethernet OAM is signaled in Ethernet packet data units.

The Maintenance end point IDs (MEP IDs) and maintenance association IDs for the switched path endpoints can be synchronized using the ETH-MCC (maintenance communication channel) transaction set once the switched path has been established.

MEPs are located at the endpoints of the Ethernet LSP. Typical configuration associated with a MEP is Maintenance Domain Name, Short Maintenance Association Name, and MA Level, MEP ID, and CCM transmission rate (when ETH-CC functionality is desired). As part of the synchronization, it is verified that the Maintenance Domain Name, Short Maintenance Association Name, MA Level, and CCM

transmission rate are the same. It is also determined that MEP IDs are unique for each MEP.

Besides the unicast CCM functionality, the PBB-TE MEPs can also offer the LBM/LBR and LMM/LMR functionalities for on-demand connectivity verification and loss measurement purposes.

[4.6](#) Protection Paths

The IEEE is currently defining protection procedures for PBB-TE [IEEE 802.1Qay]. This section will be updated when these procedures are documented.

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. Protection may be used for P2P VID/ESP-MAC DA, P2MP VID/ESP-MAC DA and P2MP VID configured modes of operation. The 'P' bit in the protection object indicates the role (working or protection) of the LSP currently being signaled.

If the initiating switch wishes to use G.8031 [[G-8031](#)] data plane protection switching coordination (vs. control plane notifications), it sets the N bit to 1 in the protection object. This must be consistently applied for all paths associated as a protection group.

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

If the terminating switch does not support G.8031, the error "Admission Control Failure/Unsupported Notification Type" is used.

[5.](#) Error conditions

The following errors have been identified as being unique to these procedures and in addition to those already defined. This will be addressed in a proper IANA considerations section in a future version of the document:

[5.1](#) Invalid ESP-VID value for PBB-TE MSTI range

The originator of the error is not configured to use the ESP-VID value in conjunction with GMPLS signaling of <ESP: VID, MAC DA > tuples. This may be any switch along the path.

[5.2](#) Invalid MAC Address

The MAC address is out of a reserved range that cannot be used by then node which is processing the address. While almost all MAC addresses are valid there are a small number of reserved MAC addresses.

[5.3](#) Invalid ERO for UPSTREAM_LABEL Object

The ERO offered has discontinuities with the identified ESP-VID/ESP-MAC DA path in the UPSTREAM_LABEL object.

[5.4](#) Invalid ERO for LABEL_SET Object

The ERO offered has discontinuities with the identified ESP-VID/ESP-MAC DA path in the LABEL_SET object.

[5.5](#) Switch is not ESP P2MP capable

This error may arise only in P2MP VID Tree allocation.

[5.6](#) Invalid ESP-VID in UPSTREAM_LABEL object

The ESP-VID in the UPSTREAM_LABEL object for the "asymmetrical ESP-VID" P2MP tree did not correspond to the ESP-VID used in previous transactions.

[6.](#) Deployment Scenarios

This technique of GMPLS controlled Ethernet switching is applicable to all deployment scenarios considered by the design team [CCAMP-ETHERNET].

[7.](#) Security Considerations

The architecture assumes that the GMPLS controlled Ethernet subnet consists of trusted devices and that the UNI ports to the domain are untrusted. Care is required to ensure untrusted access to the trusted domain does not occur. Where GMPLS is applied to the control of VLAN only, the commonly known techniques for mitigation of Ethernet DOS attacks may be required on UNI ports.

[8.](#) IANA Considerations

New values are required for signaling and error codes as indicated. This section will be completed in a later version.

[9. References](#)

[9.1 Normative References](#)

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [ARCH] Fedyk, D. Berger, L., Andersson L., "GMPLS Ethernet Label Switching Architecture and Framework", work in progress.
- [RFC3473] Berger, L. et.al., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", IETF [RFC 3473](#), January 2003.

[9.2 Informative References](#)

- [IEEE 802.1Qay] "IEEE standard for Provider Backbone Bridges Traffic Engineering", work in progress.
- [RFC4115] Aboul-Magd, O. et.al. "A Differentiated Service Two-Rate, Three-Color Marker with Efficient Handling of in-Profile Traffic", IETF [RFC 4115](#), July 2005
- [G-8031] ITU-T Draft Recommendation G.8031, Ethernet Protection Switching.
- [IEEE 802.1AB] "IEEE Standard for Local and Metropolitan Area Networks, Station and Media Access Control Connectivity Discovery".
- [IEEE 802.1ag] "IEEE Draft Standard for Connectivity Fault Management", work in progress.
- [IEEE 802.1ah] "IEEE standard for Provider Backbone Bridges", work in progress.
- [RFC4204] Lang. J. Editor, "Link Management Protocol (LMP)" [RFC4204](#), October 2005

- [MEF.6] The Metro Ethernet Forum MEF 6 (2004), "Ethernet Services Definitions - Phase I".
- [MEF.10] The Metro Ethernet Forum MEF 10 (2004), "Ethernet Services Attributes Phase 1".
- [RFC3270] Le Faucheur, F. et.al., "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services" IETF [RFC 3270](#), May 2002.
- [RFC4875] Aggarwal, R. Ed., "Extensions to RSVP-TE for Point to Multipoint TE LSPs", IETF [RFC 4875](#), May 2007
- [PATHCOMP] Farrel, A. et.al., "Path Computation Element (PCE) Architecture", work in progress.
- [RFC3985] Bryant, S., Pate, P. et al., "Pseudo Wire Emulation Edge-to Edge (PWE3) Architecture", IETF [RFC 3985](#), March 2005.
- [RFC4872] Lang et.al., "RSVP-TE Extensions in support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery", [RFC 4872](#), May 2007.
- [RFC4873] Berger, L. et.al., "MPLS Segment Recovery", [RFC 4873](#), May 2007.
- [RFC3209] Awduche et.al., "RSVP-TE: Extensions to RSVP for LSP Tunnels, IETF [RFC 3209](#), December 2001.
- [Y.1731] ITU-T Draft Recommendation Y.1731(ethoam), " OAM Functions and Mechanisms for Ethernet based Networks ", work in progress.
- [ADDRESS] Shimoto, K., Papneja, R., Rabbat, R., "Use of Addresses in Generalized Multi-Protocol Label Switching (GMPLS) Networks", work in progress.
- [CCAMP-ETHERNET] Papadimitriou, D. et.al, "A Framework for Generalized MPLS (GMPLS) Ethernet", internet draft, [draft-papadimitriou-ccamp-gmpls-ethernet-framework-00.txt](#), June 2005

10. Author's Address

Don Fedyk
Nortel Networks
600 Technology Park Drive
Billerica, MA, 01821
Email: dwfedyk@nortel.com

David Allan

Nortel Networks
3500 Carling Ave.
Ottawa, Ontario, CANADA

Fedyk et al.

Expires May 2008

[Page 18]

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

Email: dallan@nortel.com

Himanshu Shah
Ciena
35 Nagog Park,
Acton, MA 01720
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

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

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

Nurit Sprecher
Nokia Siemens Networks,
GmbH & Co. KG
C00 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

11. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has

Fedyk et al.

Expires May 2008

[Page 19]

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

12. Disclaimer of Validity

"This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

13. Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

[14](#). Acknowledgments

The authors would like to thank Dinesh Mohan, Nigel Bragg, Stephen Shew and Sandra Ballarte for their contributions to this document.

Appendix A

Rational and mechanism for PBB-TE Ethernet Forwarding

This appendix describes work currently being undertaken in the 801.1 PBB-TE [IEEE 802.1Qay] project. This information is for reference only and will be removed when 802.1Qay becomes mature. This text captures some of the original rational for changing Ethernet forwarding. The PBB-TE [IEEE 802.1Qay] document simply documents the PBB-TE data plane.

Ethernet as specified today is a complete system consisting of a data plane and a number of control plane functions. Spanning tree, data plane flooding and MAC learning combine to populate forwarding tables and produce resilient any-to-any behavior in a bridged network.

Ethernet consists of a very simple and reliable data plane that has been optimized and mass produced. By simply disabling some Ethernet control plane functionality, it is possible to employ alternative control planes and obtain different forwarding behaviors.

Customer	Provider	Provider
Bridge/	Bridge	Backbone
		Bridge


```

C-MAC/C-VID-----802.1Q -----C-MAC-CVID
      S-VID-----802.1ad-----S-VID
            B-MAC---802.1ah---B-MAC
            B-VID---802.1ah---B-VID

```

Figure 1 802.1 MAC/VLAN Hierarchy

Recent works in IETF Pseudo Wire Emulation [[RFC3985](#)] and IEEE 802 are defining a separation of Ethernet functions permitting an increasing degree of provider control. The result is that the Ethernet service to the customer appears the same, yet the provider components and behaviors have become decoupled from the customer presentation and the provider has gained control of all VID/DMAC endpoints.

One example of this is the 802.1ah work in hierarchical bridging whereby customer Ethernet frames are fully encapsulated into a provider Ethernet frame, isolating the customer VID/DMAC space from the provider VID/DMAC space. In this case, the forwarding behavior of the of the Backbone MAC in the provider's network is as per 802.1Q.

The Ethernet data plane provides protocol multiplexing via the ether type field which allows encapsulation of different protocols supporting various applications. More recently, the Carrier Ethernet effort has created provider and customer separation that enables

another level of multiplexing. This in effect creates provider MAC endpoints in the Ethernet sub-network controlled by the provider. In this appendix we concentrate on the provider solutions and therefore subsequent references to VLAN, VID and MAC refer to those under provider control, be it in the backbone layer of 802.1ah. The Customer Ethernet service is the same native Ethernet service with functions such as bridging, learning and spanning trees all functioning over the provider infrastructure.

Bridging offers a simple solution for any-to-any connectivity within a VLAN partition via the Spanning tree, flooding and MAC learning. Spanning tree provides some unnecessary capabilities for P2P services and since the Spanning tree must interconnect all MACs with the same VLAN IDs (VIDs) it consumes a scarce resource (VIDs). In this document we present that it is easier to modify Ethernet to scale engineered P2P services and this is the approach we take with

PBB-TE. (The number of usable VLANs IDs in conventional Ethernet bridging is constrained to 4094, therefore the use of VLAN only configuration for all forwarding could be limited for some applications where large number of P2P connections are required.) This is because in Ethernet, each Spanning tree is associated with one or more VLAN IDs. Also Port membership in a VLAN is configured which controls the connectivity of all MAC interfaces participating in the VLAN.

The roots for PBB-TE capability exist in the Ethernet management plane. The management of Ethernet switches provides for static configuration of Ethernet forwarding. The Ethernet Control plane allows for forwarding entries that are statically provisioned or configured. In this document we are expanding the meaning of "configured" from an Ethernet Control plane sense to mean either provisioned, or controlled by GMPLS. The connectivity aspects of Ethernet forwarding is based upon VLANs and MAC addresses. In other words the VLAN + DMAC are an Ethernet Label that can be looked up at each switch to determine the egress link (or links in the case of link aggregation).

This is a finer granularity than traditional VLAN networks since each P2P connection is independent. By provisioning MAC addresses independently of Spanning tree in a domain, both the VLAN and the VLAN/DMAC configured forwarding can be exploited. This greatly extends the scalability of what can be achieved in a pure Ethernet bridged sub network.

The global/domain wide uniqueness and semantics of MAC addresses as interface names or multicast group addresses has been preserved. (In Ethernet overlap of MAC addresses across VLANs is allowed. However for PBB-TE MAC addresses should be unique for all VLANs assigned to PBB-TE. With PBB-TE it is an operational choice if the operator uses PBT-TE labels out of the global MAC address space or the local admin space.) We then redefine the semantics associated with administration and uses of VLAN values for the case of explicit forwarding such as you get with statically configured Ethernet.

The PBB-TE is Ethernet Forwarding where configured VID + DMAC provide a forwarding table that is consistent with existing PBB and Ethernet switching. At the same time it provides domain wide labels that can be controlled by a common GMPLS control plane. This makes GMPLS control and resource management procedures ideal to create paths. The outcome is that the GMPLS control plane can be utilized

to set up the following atomic modes of connectivity:

- 1) P2P connectivity and MP2P multiplexed connectivity based on configuration of unicast MAC addresses in conjunction with a VID from a set of pre-configured VIDs.
- 2) P2MP connectivity based on configuration of multicast MAC address in conjunction with a VID from a set of pre-configured VIDs. This corresponds to (Source, Group) or (S,G) multicast.
- 3) P2MP connectivity based on configuration of VID port membership. This corresponds to (S,*) or (*,*) multicast (where * represents the extent of the VLAN Tree).
- 4) MP2MP connectivity based on configuration of VID port membership (P2MP trees in which leaves are permitted to communicate). Although, we caution that this approach poses resilience issues (discussed in [section 5](#)) and hence is not recommended.

The modes above are not completely distinct. Some modes involve combinations of P2P connections in one direction and MP connectivity in the other direction. Also, more than one mode may be combined in a single GMPLS transaction. One example is the incremental addition of a leaf to a P2MP tree with a corresponding MP2P return path (analogous to a root initiated join).

In order to realize the above connectivity modes, a partition of the VLAN IDs from traditional Ethernet needs to be established. The partition allows for a pool of Ethernet labels for manual configuration and/or for GMPLS control plane usage. The VID partition actually consists of a "configured VID/DMAC range" and "configured VID range" since in some instances the label is a VID/DMAC and sometimes the label is a VID/Multicast DMAC.

A 1. Overview of configuration of VID/DMAC tuples

Statically configured MAC and VID entries are a complete 60 bit lookup. The basic operation of an Ethernet switch is filtering on VID and forwarding on DMAC. The resulting operation is the same as performing a full 60 bit lookup (VID (12) + DMAC(48)) for P2P operations, only requiring uniqueness of the full 60 bits for forwarding to resolve correctly. This is an Ethernet domain wide label.

Complete route freedom is available for each domain wide label (60 bit VLAN/DMAC tuple) and the ability to define multiple connectivity

instances or paths per DMAC for each of the VIDs in the "configured VID/DMAC range".

The semantics of MAC addresses are preserved, and simply broaden the potential interpretations of VLAN ID from spanning tree identifier to topology instance identifier. Therefore, operation of both standard bridging and configured unicast/multicast operation is available side by side. The VID space is partitioned and a range of VIDs is allocated (say 'n' VIDs) as only significant when combined with a configured DMAC address (the aforementioned "configured VID/DMAC range" of VIDs). A VID in that range is considered as an individual connectivity instance identifier for a configured P2P path terminating at the associated DMAC address. Or in the case of P2MP, a P2MP multicast tree corresponding to the destination multicast group address. Note that this is destination based forwarding consistent with how Ethernet works today. The only thing changed is the mechanism of populating the forwarding tables.

Ethernet MAC addresses are typically globally unique since the 48 bits consists of 24 bit Organizational Unique Identifier and a 24 bit serial number. There is also a bit set aside for Multicast and for local addresses out of the OUI field. We define domain wide as within a single organization, or more strictly within a single network within an organization. For provider MAC addresses that will only be used in a domain wide sense we can define MAC addresses out of either the local space or the global space since they both have the domain wide unique property. When used in the context of GMPLS, it is useful to think of a domain wide pool of labels where switches are assigned a set of MAC addresses. These labels are assigned traffic that terminates on the respective switches.

It is also worth noting that unique identification of source in the form of the ESP-MAC SA is carried e2e in the MAC header. So although we have a 60 bit domainwide unique label, it may be shared by multiple sources and the full connection identifier for an individual P2P instance is 108 bits (ESP-MAC SA, VID and DMAC). The ESP-MAC SA is not referenced in forwarding operations but it would allow additional context for tracing or other operations at the end of the path.

For multicast group addresses, the VID/DMAC concatenated label can be distributed by the source but label assignment (as it encodes global multicast group information) requires coordination within the GMPLS controlled domain.

As mentioned earlier, this technique results in a single unique and invariant identifier, in our case a VID/DMAC label associated with the path termination or the multicast group. There can be up to

4094 labels to any one MAC address. However, practically, from Ethernet network wide aspect; there would be only a handful of VLANs allocated for PBB-TE. In addition, all 48 bits are not completely available for the MAC addresses. One way to maximize the space is to use the locally administered space. This is a large number for

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

P2P applications and even larger when shared or multiplexed forwarding is leveraged. In practice, most network scaling requirements may be met via allocation of only a small portion of the VID space, to the configured VID/DMAC range. The result is minimal impact on the number of remaining bridging VLANs that can be concurrently supported.

In order to use this unique 60 bit label, we disable the normal mechanisms by which Ethernet populates the forwarding table for the allocated range of VIDs. When a path is setup, for a specific label across a contiguous sequence of Ethernet switches, a unidirectional connection is the functional building block for an Ethernet Label Switched path (Eth-LSP).

In P2P mode a bidirectional path is composed of two unidirectional paths that are created with a single RSVP-TE session. The technique does not require the VID to be common in both directions. However, keeping in line with regular Ethernet these paths are symmetrical such that a single bidirectional connection is composed of two unidirectional paths that have common routing (i.e. traverse the same switches and links) in the network and hence share the same fate.

In P2MP mode a bidirectional path is composed of a unidirectional multicast tree and a number of P2P paths from the leaves of the tree to the root. Similarly these paths may have bandwidth and must have common routing as in the P2P case.

There are a few modifications required to standard Ethernet to make this approach robust:

1. In Standard Ethernet, discontinuities in forwarding table configuration in the path of a connection will normally result in packets being flooded as "unknown". For configured operation (e.g. PBB-TE), unknown addresses are indicative of a fault or configuration error and the flooding of these is undesirable in meshed topologies. Therefore flooding of "unknown" unicast/multicast MAC addresses must be disabled for the "configured VID/DMAC range".

2. MAC learning is not required, and although it will not interfere with management/control population of the forwarding tables, since static entries are not overridden, it appears prudent to explicitly disable MAC learning for the configured VID/DMAC and VID range.

3. Spanning tree is disabled for the allocated VID/DMAC and VID range and port blocking must be disabled to achieve complete configured route freedom. As noted earlier, it is a control plane requirement to ensure configured paths are loop free.

All three modifications described above are within the scope of acceptable configuration options defined in IEEE802.1Q specification.

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

A 2. Overview of configuration of VID port membership

Procedures almost identical to that for configuration of P2P VID/DMAC tuples can also be used for the incremental configuration of P2MP VID trees. For the replication of forwarding in this case the label is common for the multipoint destinations. The MAC field is set to multicast address and is common to the multicast community. The VID is a distinguisher common to the multicast community. The signaling procedures are as per that for [[RFC4875](#)].

Since VID translation is relatively new and is not a ubiquitously deployed capability, we consider a VID to be a domain global value. Therefore, the VID value to be used by the originating switch may be assigned by management and nominally is required to be invariant across the network. The ability to indicate permissibility of translation will be addressed in a future version of the document.

A procedure known as "asymmetrical VID" may be employed to constrain connectivity (root to leaves, and leaves to root only) when switches also support shared VLAN learning (or SVL). This would be consistent with the root as a point of failure.

A 3. OAM Aspects

Robustness is enhanced with the addition of data plane OAM to provide both fault and performance management.

For the configured VID/DMAC unicast mode of behavior, the hardware performs unicast packet forwarding of known MAC addresses exactly as

Ethernet currently operates. The OAM currently defined, [802.1ag and Y.1731] can also be reused without modification of the protocols. However currently if the VID for PBB-TE is different in each direction some modification of the OAM may be required.

An additional benefit of domain wide path identifiers, for data plane forwarding, is the tight coupling of the 60 bit unique connection ID (VID/DMAC) and the associated OAM packets. It is a simple matter to determine a broken path or misdirected packet since the unique connection ID cannot be altered on the Eth-LSP. This is in fact one of the most powerful and unique aspects of the domain wide label for any type of rapid diagnosis of the data plane faults. It is also independent of the control plane so it works equally well for provisioned or GMPLS controlled paths.

Bidirectional transactions (e.g. ETH-LB) and reverse direction transactions MAY have a different VID for each direction. PBB-TE is specifying this aspect of CFM.

For configured multicast VID/DMAC mode, the current versions of 802.1ag and Y.1731] make no representation as to how PDUs which are not using unicast addresses or which use OAM reserved multicast

Internet Draft [draft-fedyk-gmpls-ethernet-pbt-te-02.txt](#)

addresses are handled. Therefore this specification makes no representation as to whether such trees can be instrumented.

When configured VID mode of operation is used PBB-TE can be forced to use the same VID in both directions, emulating the current Ethernet data plane and the OAM functions as defined in the current versions of 802.1ag and Y.1731 can be used with no restriction.

A 4. QOS Aspects

Ethernet VLAN tags include priority tagging in the form of the 802.1p priority bits. When combined with configuration of the paths via management or control plane, priority tagging produces the Ethernet equivalent of an MPLS-TE E-LSPs [RFC3270]. Priority tagged Ethernet PDUs self-identify the required queuing discipline independent of the configured connectivity.

It should be noted that the consequence of this is that there is a common COS model across the different modes of configured operation specified in this document.

The actual QOS objects required for signaling will be in a future version of this memo.

A 5. Resiliency Aspects

A 5.1. E2E Path protection

One plus One(1+1) protection is a primary LSP with a disjoint dedicated back up LSP. One for one (1:1) protection is a primary LSP with a disjoint backup LSP that may share resources with other LSPs. One plus One and One for One Automatic Protection Switching strategies are supported. Such schemes offer:

- 1) Engineered disjoint protection paths that can protect both directions of traffic.
- 2) Fast switchover due to tunable OAM mechanisms.
- 3) Revertive path capability when primary paths are restored.
- 4) Option for redialing paths under failure.

Specific procedures for establishment of protection paths and associating paths into "protection groups" are TBD.

Note that E2E path protection is able to respond to failures with a number of configurable intervals. The loss of CCM OAM frames in the data plane can trigger paths to switch. In the case of CCM OAM frames, the detection time is typically 3.5 times the CCM interval plus the propagation delay from the fault.