

Internet Draft
Category: Informational
Expiration Date: May 9, 2008

Don Fedyk, Nortel
Lou Berger, LabN
Loa Andersson, Acreo AB

November 9, 2007

GMPLS Ethernet Label Switching Architecture and Framework

[draft-gmpls-ethernet-arch-00.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/lid-abstracts.html>

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

This Internet-Draft will expire on May 9, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

There has been significant recent work in increasing the capabilities of Ethernet switches. As a consequence, the role of Ethernet is rapidly expanding into "transport networks" that previously were the domain of other technologies such as SONET/SDH TDM and ATM. This document defines an architecture and framework for a GMPLS based control plane for Ethernet in this "transport network" capacity. GMPLS has already been specified for similar technologies. Some additional extensions to the GMPLS control plane are needed and this document provides a framework for these extensions.

Internet-Draft

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

November 9, 2007

Contents

1	Introduction	3
2	Background	5
2.1	Ethernet Switching	5
2.2	Operations, Administration, and Maintenance (OAM)	7
2.3	Terminology	8
2.3.1	Concepts	8
2.3.2	Acronyms	9
2.4	Ethernet and MPLS similarities and differences	10
3	Framework	10
4	GMPLS Routing and Addressing Model	12
4.1	GMPLS Routing	13
4.2	Control Plane Network	13
5	P2P Signaling	13
6	Link Management	14
7	Path Computation and Selection	15
8	Multiple Domains	16
9	Security Considerations	16
10	IANA Considerations	16
11	References	16
11.1	Normative References	16
11.2	Informative References	16
12	Acknowledgments	18
13	Author's Addresses	18
14	Full Copyright Statement	18
15	Intellectual Property	19

Internet-Draft

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

November 9, 2007

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 is the initial draft of this document.

[1](#). Introduction

There has been significant recent work in increasing the capabilities of Ethernet switches. As a consequence, the role of Ethernet is rapidly expanding into "transport networks" that previously were the domain of other technologies such as SONET/SDH TDM and ATM. The evolution and development of Ethernet capabilities in these areas is a very active and ongoing process.

Multiple organizations have been active in extending Ethernet technology. This activity has taken place in the IEEE 802.1 Working Group, the ITU and the MEF. These groups have been focusing on Ethernet forwarding, Ethernet management plane extensions and the Ethernet Spanning Tree Control Plane, but not an explicitly routed, constraint based control plane.

In the forwarding plane context, extensions have been, or are being, defined to support different Ethernet forwarding models, protection modes and service interfaces. Examples of such extensions include [[802.1ah](#)], [[802.1Qay](#)], [[G.8011](#)] and [[MEF.6](#)]. The provider extensions allow for greater flexibility in the forwarding plane. In the 802.1Qay case the extensions allow for a departure from forwarding based on Ethernet spanning tree. The greater flexibility in forwarding is achieved through the addition of a "provider" address

space.

This document is a framework for GMPLS Ethernet Label switching (GELS). It will be followed by technology specific documents. GELS will require more than one switching type, and the GMPLS procedures that will need to be changed is dependent on switching, and thus will be covered in the technology specific documents.

In the new provider bridge model developed in the IEEE802.1ad-project and amended to the IEEE802.1Q standard [[802.1Q](#)], an extra VLAN identifier (VID) is added. This VLAN is referred to as the Service VID, (S-VID and is carried in a Service TAG (S-TAG). In provider backbone bridges (PBB) [[802.1ah](#)] a backbone VID (B-VID) and B-MAC

header with a Service Instance (I-TAG) encapsulates a customer Ethernet frame or a service Ethernet frame. An example of Ethernet protection extensions can be found in [[G.8031](#)]. In the IEEE802.1Q standard the terms Provider Backbone Bridges (PBB) and Provider Backbone Bridged Network (PBBN) is used in the context of these extensions.

Ethernet operations, administration, and maintenance (OAM) is another important area that is being extended to enable provider Ethernet services. Related extensions can be found in [[802.1ag](#)] and [[Y.1731](#)].

The Ethernet forwarding and management plane extensions explicitly allow for the disabling of standard Ethernet spanning tree but do not define an explicitly routed, constraint based control plane. The IEEE802.1, in [[802.1Qay](#)], works on an new amendment that explicitly allows for traffic engineering of Ethernet forwarding paths.

The IETF chartered the GMPLS work to specify a common control plane for physical path and core tunneling technologies for the Internet and telecommunication service providers. The GMPLS architecture is specified in [RFC3945](#) [[RFC3945](#)]. The protocols specified for GMPLS have been used to control "Transport Networks", e.g. Optical and TDM networks.

This document provides a framework for use of GMPLS to control "transport" Ethernet. The GMPLS architecture already handles a number of transport technologies but Ethernet adds a few new constraints that must be documented. Some additional extensions to the GMPLS

control plane are needed and this document provides a framework for these extensions.

This document introduces and explains the concept of an Ethernet Label Switched Path (Eth-LSP). The data plane aspects of Eth-LSPs are outside the scope of this document and IETF activities.

The intent of this document is to reuse and align with as much of the GMPLS protocols as possible. For example reusing the IP control plane addressing allows the existing signaling, routing, LMP and path computation to be used as specified. Additions are made only to accommodate features of Ethernet that are not already supported by GMPLS. The GMPLS protocols support a set of tools for hierarchical LSPs as well as contiguous LSPs. GMPLS specific protocol mechanisms support a variety of networks from peer to peer to UNIs and NNIs.

[2.](#) Background

This section provides background to the types of switching and services that are supported within the defined framework. The former is particularly important as it identifies the switching functions that GMPLS will need to represent and control. The intent is for this document to allow for all standard forms of Ethernet switching and services.

The material presented in this section is based on the on-going work taking place in the IEEE 802.1 Working Group, the ITU and the MEF. This section references and, to some degree, summarizes that work. This section is not a replacement for, or an authoritative description of that work.

[2.1.](#) Ethernet Switching

In Ethernet switching terminology, the bridge relay is responsible for forwarding and replicating the frames. Bridge relays forward frames based on the header fields: Virtual Local Area Network (VLAN)

Identifiers (VID) and Destination Media Access Control (DMAC) address. PBB [802.1ah] has also introduced a Service Instance tag (I-TAG). Across all the Ethernet extensions (already referenced in the Introduction), multiple forwarding functions, or service interfaces, have been defined using the combination of VIDs, DMACs, and I-TAGs. PBB [802.1ah] provides a breakdown of the different types of Ethernet switching services. Figure 1 reproduces this breakdown.




Figure 1: Ethernet Switching Service Types

The types are defined in Clause 25 of [802.1ah], and are consistent with the definitions of Ethernet services supported in [G.8011] and [MEF.6]. To summarize the definitions:

- o Port based
This is a frame based service that supports specific frame types, no Service VLAN tagging, with MAC address based switching.
- o S-tagged
There are multiple Service VLAN tag (S-tag) aware services, including:
 - + one-to-one
In this service, each VLAN identifier (VID) is mapped into a different service.
 - + Bundled
Bundled S-tagged service supports the mapping of multiple VIDs into a single service and include:
 - * many-to-one
In this frame based service, multiple VIDs are mapped into the same service.
 - * all-to-one
In this frame based service, all VIDs are mapped into the same service.

- transparent
This is a special case, all frames are mapped from a single incoming port to a single destination Ethernet port.
- o I-tagged
The edge of a PBBN consists of a combined backbone relay (B-component relay) and service instance relay (I-component relay).

An I-Tag contains a service identifier (24 bit I-SID) and priority markings as well as some other flags. An I-Tagged service is typically between the edges of the PBBN and terminated at each edge on an I-component that faces a customer port so the service is often not visible except at the edges. However since the I-component relay involves a distinct relay it is possible to have a visible I-Tagged Service by separating the I component relay from the B-component relay. Two examples where it makes sense to do this are: an I-Tagged service between two PBBNs and as an attachment to a customer's Provider Instance Port.

In general, the different types determine which of the Ethernet header fields are used in the forwarding/switching function, e.g. VID only or VID and DMACs. The types may also require the use of additional Ethernet headers or fields. Services defined for UNIs tend to use the headers on a hop-by-hop basis.

In most cases for bridging, the header fields cannot be changed hop-by-hop, but some translations of VID field values are permitted typically at the edges. Again, while not specifically described in 802.1ah, the Ethernet services being defined in the context of [MEF.6] and [G.8011] also fall into the types defined in Figure 1.

Across all service types, the Ethernet data plane is bi-directional congruent. This means that the forward and reverse paths share the exact same set of nodes, ports and bi-directional links. This property is fundamental. The 802.1 group has defined maintained this bi-directional congruent property in the definition of Connectivity Fault Management (CFM) which is part of the overall Operations Administration and Management (OAM) capability.

2.2. Operations, Administration, and Maintenance (OAM)

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

For the Eth LSP 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.

OAM relies on domain wide path identifiers, for data plane forwarding, utilizing the 60 bit unique connection ID (VID/DMAC). Determining a broken path or misdirected packet in this case relies on the connection ID not being altered on the Eth-LSP. These identifiers are independent of the control plane so it works equally well for provisioned or GMPLS controlled paths.

Ethernet OAM currently consists of:

Defined in both [802.1ag & Y.1731]:

- CCM/RDI: Connectivity Check, Remote Defect Indication
- LBM/LBR: Loopback Message, Loopback Reply
- LTM/LTR: Link trace Message, Link trace Reply
- VSM/VSR: Vendor-specific extensions Message/Reply

Additionally defined in [[Y.1731](#)]:

- AIS: Alarm Indication Signal
- LCK: Locked Signal
- TST: Test
- LMM/LMR: Loss Measurement Message/Reply
- DM/DMM/DMR: Delay Measurement
- EXM/EXR: Experimental
- APS, MCC: Automatic Protection Switching, Maintenance Communication Channel
- Placeholders for ITU other standards

With some Eth-LSP label formats bidirectional transactions (e.g. LBM/LBR) and reverse direction transactions MAY have a different VID for each direction. Currently Y.1731 & 802.1ag makes no representations with respect to this but work is underway to address this in PBB-TE [[802.1Qay](#)].

[2.3](#). Terminology

[2.3.1](#). Concepts

The following are basic Ethernet and GMPLS terms:

o Asymmetric Bandwidth

This term refers to the property of a Bi-directional LSP may have differing bandwidth allocation in each direction.

Internet-Draft

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

November 9, 2007

- o Bi-directional Congruent LSP

This term refers to the property that an LSP shared the same nodes, ports and links. Ethernet data planes are normally bi-directional congruent.

- o Shared forwarding

Shared forwarding is a property of a data path where a single forwarding entry (VID + DMAC) may be used for frames from multiple sources (SMAC). Shared forwarding does not change any data plane behavior it saves forwarding information base (FIB) entries only. From all other aspects it behaves as if there were multiple FIB entries.

- o In-band GMPLS Signaling

In-band GMPLS Signaling is IP based signaling on the native Ethernet links encapsulated by a single hop Ethernet header. Logical links that use a dedicated VID on the same physical links would be considered In-band signaling.

- o Out-of-band GMPLS Signaling is IP based signaling between Ethernet switches that uses some other links other than the Ethernet data plane links. Out of band signaling typically shares a different fate from the data links.

- o Contiguous Eth-LSP is an Eth-LSP that maps one to one with and LSP at a domain boundary. Stitched LSP are contiguous LSPs.

- o Hierarchical Eth-LSPs are Eth-LSPs that are encapsulated and tunneled either individually or bundled with other LSPs through a domain.

[2.3.2.](#) Acronyms

The following acronyms are used in this document:

CFM	Connectivity Fault Management
DMAC	Destination MAC Address
CCM	Continuity Check Message
Eth-LSP	Ethernet Label Switched Path

I-SIDService Identifier	
LMP	Link Management Protocol
MAC	Media Access Control
MP2MP	Multipoint to multipoint
NMS	Network Management System

OAM	Operations, Administration and Maintenance
PBB	Provider Backbone Bridges [802.1ah]
PBB-TE	Provider Backbone Bridges Traffic Engineering [802.1Qay]
P2P	Point to Point
P2MP	Point to Multipoint
QoS	Quality of Service
SMAC	Source MAC Address
S-TAG	A service TAG defined in the 802.1 Standard [802.1Q]
TE	Traffic Engineering
TAG	An Ethernet short form for a TAG Header
TAG Header	An extension to an Ethernet frame carrying priority and other information.
TSpec	Traffic specification
VID	VLAN Identifier
VLAN	Virtual LAN

[2.4.](#) Ethernet and MPLS similarities and differences

Ethernet is similar to MPLS in that there is a default payload type. In MPLS the default payload is either another MPLS label or an IP packet. The IP packet may carry any type of service IP carries. Ethernet assumes an Ethernet frame as the default payload. The actual service can be anything that Ethernet carries.

In MPLS pseudo wires where other types of payloads are used natively the payload may be identified implicitly or explicitly by using a control word removing the need for the IP header.

Similarly, in Ethernet the option to carry other payloads by using either implicit or explicit means is being discussed.

Ethernet bridging is different from MPLS in that while the switching decision is taken on whatever is defined as the Ethernet label, that

label is usually not swapped at each hop.

3. Framework

As defined in the (GMPLS) Architecture [[RFC3945](#)], the GMPLS control plane can be applied to a technology by controlling the data plane and switching characteristics of that technology. The architecture includes a clear separation between a control plane and a data plane. Control plane and data plane separation allows the GMPLS control plane to remain architecturally and functionally unchanged while controlling different technologies. The architecture also requires

IP connectivity for the control plane to exchange information, but does not otherwise require an IP data plane.

All aspects of GMPLS, i.e., addressing, signaling, routing and link management, may be applied to Ethernet switching. GMPLS can provide control for traffic engineered and protected Ethernet service paths. This document defines the term "Eth-LSP" to refer to Ethernet service paths that are controlled via GMPLS. As is the case with all GMPLS controlled services, Eth-LSPs can leverage common traffic engineering attributes such as:

- bandwidth profile;
- priority level;
- preemption characteristics;
- protection/resiliency capability;
- routing policy, such as an explicit route;
- bi-directional service;
- end-to-end and segment protection;
- hierarchy

The bandwidth profile may be used, to set committed information rate and peak information rate, and policies based on either under-subscription or over-subscription. Services covered by this framework MUST have a TSpec that follows the Ethernet Traffic parameters defined in [[ETH-TSPEC](#)].

In applying GMPLS to Ethernet, GMPLS will be extended to work with the Ethernet data plane and switching functions. The definition of GMPLS support for Ethernet is multi-faceted due to the different

forwarding/switching functions inherent in the different service types discussed in [Section 2.1](#). In general, the header fields used in the forwarding/switching function, e.g. VID and DMAC, can be characterized as a data plane label. In some circumstances these fields will be constant along the path of the Eth-LSP, and in others they may vary hop-by-hop or at certain interfaces only along the path. In the case where the "labels" must be forwarded unchanged, there are a few constraints on the label allocation that are similar to some other technologies such as lambda labels.

The general characteristics of the IEEE 802.1Q [\[802.1Q\]](#) data plane are left unchanged. The VID is used as a "filter" pointing to a particular forwarding table, and if the DMAC is found in that forwarding table the forwarding decision is taken based on the DMAC. When forwarding using an Ethernet spanning tree, if the DMAC is not found the frame is broadcast over all outgoing interfaces for which that VID is defined. This valid MAC checking and broadcast supports Ethernet learning. The amendment to IEEE802.1Q that is specified under IEEE802.1Qay allows for turning off learning and hence this

broadcast mechanism. A special case is when a VID is defined for only two ports on one bridge, in that case all frames with that VID received over one of these ports are forward over the over port.

This document does not define any specific format for an Eth-LSP label. Rather, it is expected that service specific documents will define any signaling and routing extensions needed to support that specific Ethernet service. Depending on the requirements of a service, it may be necessary to define multiple GMPLS extensions, e.g., label formats, switching types, and Traffic Engineering (TE) routing extensions. It is expected that all such extensions will be consistent with this document. It is expected that there will be no case where an Eth-LSP will be signaled, or an Eth-LSP supporting interface will be represented, using the L2SC switching type/capability. A new switching type/capability is required to avoid any potential current usage of the L2SC switching type/capability in support of Ethernet.

For discussion purposes, we decompose the problem of applying GMPLS into the functions of Routing, Signaling, Link Management and Path Selection. It is possible to use some functions of GMPLS alone or in partial combinations. In most cases using all functions of GMPLS

leads to less operational overhead than partial combinations.

4. GMPLS Routing and Addressing Model

The GMPLS Routing and Addressing Model is not modified by this document. GMPLS control for Eth-LSPs uses the Routing and Addressing Model described in [\[RFC3945\]](#). Most notably this includes the use of IP addresses to identify interfaces and LSP end-points. It also includes support for both numbered and unnumbered interfaces.

In the case where another address family or type of identifier is required to support an Ethernet service, extensions may be defined to provide mapping to an IP address. Extensions to support non-IP based LSP identification in signaling, i.e., replacement of the IP address in the RSVP SESSION or SENDER_TSPEC objects, are not permitted under this framework.

4.1. GMPLS Routing

GMPLS routing [\[RFC4202\]](#) is IP routing with the opaque TLV extensions for the purpose of distributing GMPLS related TE (router and link) information. As is always the case with GMPLS, TE information is populated with TE resources coordinated with LMP or from configuration if LMP is not available. The bandwidth resources of the links are tracked as Eth-LSPs are set up. Interfaces supporting the switching of Eth-LSPs are identified using the appropriate Interface Switching Capabilities. As mentioned in [Section 2](#), it will be necessary to define one or more new Interface Switching Capabilities to support Eth-LSPs. The L2SC Interface Switching Capabilities MUST NOT be used to represent interfaces capable of supporting Eth-LSPs. Interface Switching Capability specific TE information may be defined as needed to support the requirements of a specific Ethernet

Switching Service Type.

GMPLS Routing is an optional piece but it is highly valuable in maintaining topology and distributing the TE database for path management and dynamic path computation.

[4.2.](#) Control Plane Network

In order for a GMPLS control plane to operate, an IP network of sufficient capacity to handle the information exchange between the GMPLS routing and signaling protocols is necessary.

One way to implement this is with an IGP that views each switch as a terminated IP adjacency. In other words, IP traffic and a simple routing table are available for the control plane but there is no requirement for a high performance IP data plane.

This IP connectivity can be provided as a separate independent network (out of band) or integrated with the Ethernet switches (in-band).

[5.](#) P2P Signaling

GMPLS signaling, see [[RFC3471](#)], is well suited to the control of Eth-LSPs and Ethernet switches. Signaling enables the ability to dynamically establish a path from one ingress or egress node. The signaled path may be completely static and not change for the duration of its lifetime. However, signaling also has the capability to dynamically adjust the path in a coordinated fashion after the path has been established. The range of signaling options from static to dynamic are under operator control. Standardized signaling also

improves multi-vendor interoperability over simple management.

GMPLS signaling supports the establishment and control of bidirectional and unidirectional data paths. Ethernet is bi-directional by nature and the CFM has been built to leverage this. Prior to CFM the emulation of a physical wire and the learning requirements also mandated bi-direction connections. Given this, Eth-LSPs MUST always use paths that share the same routes and fates. Eth-

LSPs may be either P2P or P2MP (see [[RFC4875](#)]). GMPLS signaling also allows for full and partial LSP protection; see [[RFC4872](#)] and [[RFC4873](#)].

Note that standard GMPLS does not support different bandwidth in each direction of a bidirectional LSP. See [[GMPLS-ASYM](#)] if asymmetric bandwidth bidirectional LSPs are required.

6. Link Management

Link discovery has been specified for Ethernet in [[IEEE802.1AB](#)]. However the 802.1AB capability is an optional feature and is not necessarily operating before the Link is operational it primarily supports the management plane. The benefits of running link discovery in large systems are significant. Link discovery may reduce configuration and reduce the possibility of undetected errors in configuration as well as exposing misconnections.

In the GMPLS context, LMP [[RFC4204](#)] has been defined to support link management and discovery features. LMP also supports the creation of unnumbered interfaces can be automated. If LMP is not used there is an additional provisioning requirement to add GMPLS link identifiers. For large-scale implementations LMP would be beneficial. LMP also has Fault Management capabilities that overlap with [[IEEE802.1ag](#)] and [[Y.1731](#)]. It is recommended that LMP not be used for Fault management and instead the native Ethernet methods be used.

LMP and 802.1AB are relatively independent. The LMP capability should be sufficient to remove the need for 802.1AB but 802.1 AB can be run in parallel or independently if desired. Figure 2 provides possible ways of using LMP, 802.1AB and 802.1ag in combination.

Figure 2 illustrates the functional relationship of link management and OAM schemes. It is intended that LMP would use functions of link property correlation but that Ethernet mechanisms for OAM such as CFM, link trace etc would be used for fault management and fault trace.

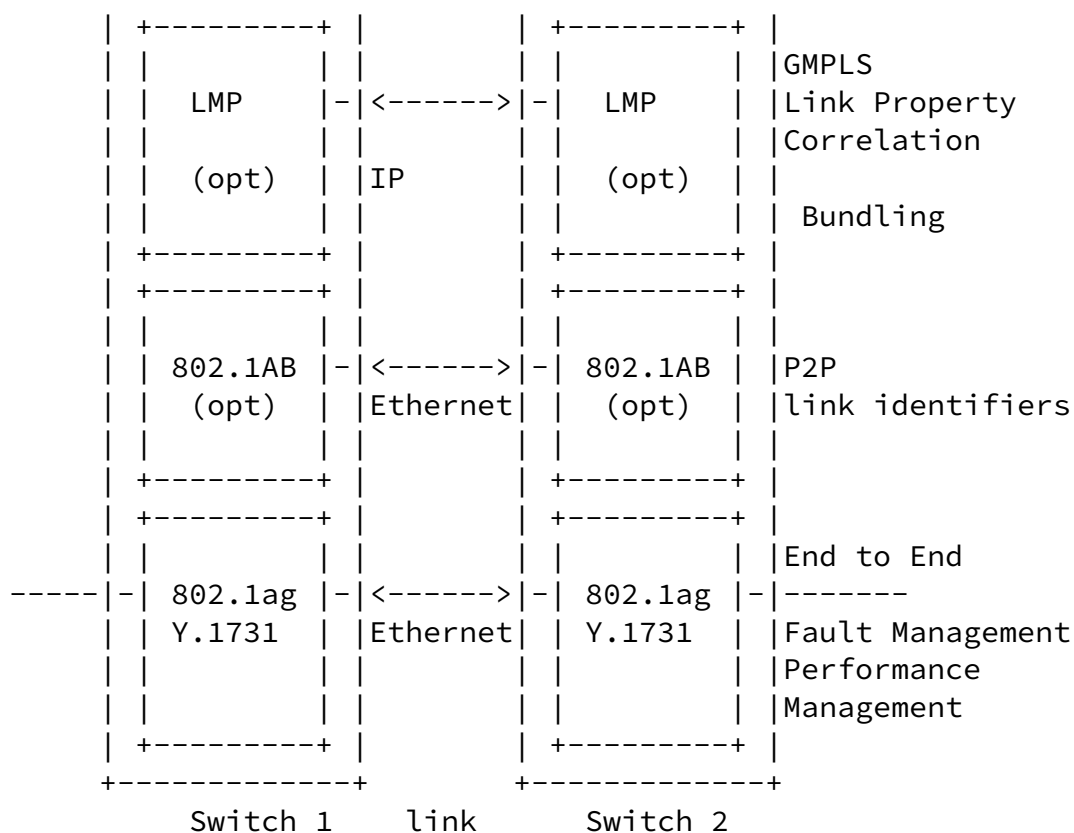


Figure 2: Logical Link Management Options

7. Path Computation and Selection

GMPLS does not specify a specific method for selecting paths or supporting path computation. GMPLS allows for a wide ranges of possibilities supported from very simple path computation to very elaborate path coordination where a large number of coordinated paths are required. The path computation could take the form of paths being computed either on a management station with local computation for rerouting or more sophisticated path computation servers.

Eth-LSPs may be supported using any path selection or computation mechanism. As is the case with any GMPLS path selection function, and common to all path selection mechanisms, the path selection process should take into consideration Switching Capabilities and Encoding advertised for a particular interface. Eth-LSPs may also make use of the emerging path computation and selection work; see [[PATH-COMP](#)]

[8. Multiple Domains](#)

This document will support the signaling of Ethernet parameters across multiple domains supporting both contiguous Eth-LSP and Hierarchical Ethernet LSPs. The intention is to support the GMPLS tools of hierarchy supporting Peer to Peer models, UNIs and NNIs.

More detail will be added to the section in a later revision.

[9. Security Considerations](#)

The architecture for GMPLS controlled Ethernet assumes that the network consists of trusted devices, but does not require that the ports over which a UNI is defined is trusted, nor does equipment connected to these ports need to be trusted. Access to the trusted domain SHALL only occur through the protocols defined in the UNI or NNI or through protected management interfaces. 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.

[10. IANA Considerations](#)

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

[11. References](#)

[11.1. Normative References](#)

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3471] Berger, L. (editor), "Generalized MPLS Signaling Functional Description", January 2003, [RFC3471](#).
- [RFC4202] Kompella, K., Rekhter, Y., "Routing Extensions in Support of Generalized MPLS", [RFC 4202](#), October 2005

[11.2. Informative References](#)

- [G.8031] ITU-T Draft Recommendation G.8031, Ethernet Protection Switching.

Internet-Draft

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

November 9, 2007

- [G.8011] ITU-T Draft Recommendation G. 8011, Ethernet over Transport - Ethernet services framework.
- [RFC3945] E. Mannie, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", [RFC 3495](#).
- [IEEE802.1AB] "IEEE Standard for Local and Metropolitan Area Networks, Station and Media Access Control Connectivity Discovery".
- [IEEE802.1ag] "IEEE Draft Standard for Connectivity Fault Management", work in progress.
- [802.1ah] "IEEE standard for Provider Backbone Bridges", work in progress.
- [802.1Qay] "IEEE standard for Provider Backbone Bridge Traffic Engineering", work in progress.
- [802.1Q] "IEEE standard for Virtual Bridged Local Area Networks 802.1Q-2005", May 19, 2006
- [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".
- [RFC4875] Aggarwal, R. Ed., "Extensions to RSVP-TE for Point to Multipoint TE LSPs", IETF [RFC 4875](#), May 2007
- [PATH-COMP] Farrel, A. et.al., "Path Computation Element (PCE) Architecture", work in progress.
- [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.

[Y.1731] ITU-T Draft Recommendation Y.1731(ethoam), " OAM Functions and Mechanisms for Ethernet based Networks ", work in progress.

Fedyk, et. al.

Informational

[Page 17]

Internet-Draft

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

November 9, 2007

[GMPLS-ASYM] Berger, L. et al., "GMPLS Asymmetric Bandwidth Bidirectional LSPs", work in progress.

[ETH-TSPEC] Papadimitriou, D., "Ethernet Traffic Parameters", work in progress.

[12.](#) Acknowledgments

There were many people involved in the initiation of this work prior to this document. The GELS framework draft and the PBB-TE extensions drafts were two drafts that helped shape and justify this work. We acknowledge the work of these authors of these initial drafts: Dimitri Papadimitriou, Nurit Sprecher, Jaihyung Cho, Dave Allan, Peter Busschbach, Attila Takacs, Thomas Eriksson, Diego Caviglia, Himanshu Shah, Greg Sunderwood, Alan McGuire, Nabil Bitar.

George Swallow contributed significantly to this document.

[13.](#) Author's Addresses

Don Fedyk
Nortel Networks
600 Technology Park Drive
Billerica, MA, 01821
Phone: +1-978-288-3041
Email: dwfedyk@nortel.com

Lou Berger
LabN Consulting, L.L.C.
Phone: +1-301-468-9228

Email: lberger@labn.net

Loa Andersson
Acreo, AB
Phone:+46 8 632 77 14
Email: loa@pi.se

14. Full 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.

This document and the information contained herein are provided on an

Fedyk, et. al.

Informational

[Page 18]

Internet-Draft

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

November 9, 2007

"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.

15. Intellectual Property

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 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.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).