Internet Engineering Task Force (IETF)
Request for Comments: 5828
Category: Informational
ISSN: 2070-1721

D. Fedyk
Alcatel-Lucent
L. Berger
LabN
L. Andersson
Ericsson

March 2010

Generalized Multiprotocol Label Switching (GMPLS) Ethernet Label Switching Architecture and Framework

Abstract

There has been significant recent work in increasing the capabilities of Ethernet switches and Ethernet forwarding models. As a consequence, the role of Ethernet is rapidly expanding into "transport networks" that previously were the domain of other technologies such as Synchronous Optical Network (SONET) / Synchronous Digital Hierarchy (SDH), Time-Division Multiplexing (TDM), and Asynchronous Transfer Mode (ATM). This document defines an architecture and framework for a Generalized-MPLS-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.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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/rfc5828.

Table of Contents

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 to support transport networks. This activity has taken place in the Institute of Electrical and Electronics Engineers (IEEE) 802.1 Working Group, the International Telecommunication Union - Telecommunication Standardization Sector (ITU-T) and the Metro Ethernet Forum (MEF). These groups have been focusing on Ethernet forwarding, Ethernet management plane extensions, and the Ethernet Spanning Tree Control Plane, but not on an explicitly routed, constraint-based control plane.

In the forwarding-plane context, extensions have been, or are being, defined to support different transport Ethernet forwarding models, protection modes, and service interfaces. Examples of such extensions include [802.1ah], [802.1Qay], [G.8011], and [MEF.6]. These extensions allow for greater flexibility in the Ethernet forwarding plane and, in some cases, the extensions allow for a departure from forwarding based on a spanning tree. For example, in the [802.1ah] case, greater flexibility in forwarding is achieved through the addition of a "provider" address space. [802.1Qay] supports the use of provisioning systems and network control protocols that explicitly select traffic-engineered paths.

This document provides a framework for GMPLS Ethernet Label Switching (GELS). GELS will likely require more than one switching type to support the different models, and as the GMPLS procedures that will need to be extended are dependent on switching type, these will be covered in the technology-specific documents.

In the provider bridge model developed in the IEEE 802.1ad project and amended to the IEEE 802.1Q standard [802.1Q], an extra Virtual Local Area Network (VLAN) identifier (VID) is added. This VID is referred to as the Service VID (S-VID) and is carried in a Service TAG (S-TAG). In Provider Backbone Bridges (PBBs) [802.1ah], a Backbone VID (B-VID) and B-MAC header with a service instance (I-TAG) encapsulate a customer Ethernet frame or a service Ethernet frame.

In the IEEE 802.1Q standard, the terms Provider Backbone Bridges (PBBs) and Provider Backbone Bridged Network (PBBN) are used in the context of these extensions.

An example of Ethernet protection extensions can be found in [G.8031]. 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].

An Ethernet-based service model is being defined within the context of the MEF and ITU-T. [MEF.6] and [G.8011] provide parallel frameworks for defining network-oriented characteristics of Ethernet services in transport networks. These framework documents discuss general Ethernet connection characteristics, Ethernet User-Network Interfaces (UNIs), and Ethernet Network-Network Interfaces (NNIs). [G.8011.1] defines the Ethernet Private Line (EPL) service, and [G.8011.2] defines the Ethernet Virtual Private Line (EVPL) service. [MEF.6] covers both service types. These activities are consistent with the types of Ethernet switching defined in [802.1ah].

The Ethernet forwarding-plane and management-plane extensions allow for the disabling of standard Spanning Tree Protocols but do not define an explicitly routed, constraint-based control plane. For example, [802.1Qay] is an amendment to IEEE 802.1Q that explicitly allows for traffic engineering of Ethernet forwarding paths.

The IETF’s GMPLS work provides a common control plane for different data-plane technologies for Internet and telecommunication service providers. The GMPLS architecture is specified in RFC 3945 [RFC3945]. The protocols specified for GMPLS can be used to control "Transport Network" technologies, e.g., optical and TDM networks. GMPLS can also be used for packet and Layer 2 Switching (frame/cell-based networks).

This document provides a framework for the use of GMPLS to control "transport" Ethernet Label Switched Paths (Eth-LSPs). Transport Ethernet adds new constraints that require it to be distinguished from the previously specified technologies for GMPLS. Some additional extensions to the GMPLS control plane are needed, and this document provides a framework for these extensions. All extensions to support Eth-LSPs will build on the GMPLS architecture and related specifications.

This document introduces and explains GMPLS control plane use for transport Ethernet and the concept of the 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 be aligned with as much of the GMPLS protocols as possible. For example, reusing the IP control-plane addressing allows existing signaling, routing, Link

Management Protocol (LMP), and path computation to be used as specified. The GMPLS protocols support hierarchical LSPs as well as contiguous LSPs. Also, GMPLS protocol mechanisms support a variety of network reference points from UNIs to NNIs. Additions to existing GMPLS capabilities will only be made to accommodate features unique to transport Ethernet.

1.1. Terminology

1.1.1. Concepts

The following are basic Ethernet and GMPLS terms:

- Asymmetric Bandwidth

  This term refers to a property of a bidirectional service instance that has differing bandwidth allocation in each direction.

- Bidirectional congruent LSP

  This term refers to the property of a bidirectional LSP that uses only the same nodes, ports, and links in both directions. Ethernet data planes are normally bidirectional congruent (sometimes known as reverse path congruent).

- Contiguous Eth-LSP

  A contiguous Eth-LSP is an end-to-end Eth-LSP that is formed from multiple Eth-LSPs, each of which is operating within a VLAN and is mapped one-to-one at the VLAN boundaries. Stitched LSPs form contiguous LSPs.

- Eth-LSP

  This term refers to Ethernet Label Switched Paths that may be controlled via GMPLS.

- Hierarchical Eth-LSP

  Hierarchical Eth-LSPs create a hierarchy of Eth-LSPs.

- In-band GMPLS signaling

  In-band GMPLS signaling is composed of IP-based control messages that are sent 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.

- Out-of-band GMPLS signaling

  Out-of-band GMPLS signaling is composed of IP-based control messages that are sent between Ethernet switches over links other than the links used by the Ethernet data plane. Out-of-band signaling typically shares a different fate from the data links.

- Point-to-point (P2P) Traffic Engineering (TE) service instance

  A TE service instance made up of a single bidirectional P2P or two P2P unidirectional Eth-LSPs.

- Point-to-multipoint (P2MP) Traffic Engineering (TE) service instance

  A TE service instance supported by a set of LSPs that comprises one P2MP LSP from a root to n leaves, plus a bidirectional congruent point-to-point (P2P) LSP from each of the leaves to the root.

- Shared forwarding

  Shared forwarding is a property of a data path where a single forwarding entry (VID + Destination MAC address) may be used for frames from multiple sources (Source MAC addresses). Shared forwarding does not change any data-plane behavior. Shared forwarding saves forwarding database (FDB) entries only. Shared forwarding offers similar benefits to merging in the data plane. However, in shared forwarding, the Ethernet data packets are unchanged. With shared forwarding, dedicated control-plane states for all Eth-LSPs are maintained regardless of shared forwarding entries.

1.1.2. Abbreviations and Acronyms

The following abbreviations and acronyms are used in this document:

CCM	Continuity Check Message
CFM	Connectivity Fault Management
DMAC	Destination MAC Address
Eth-LSP	Ethernet Label Switched Path
I-SID	Backbone Service Identifier carried in the I-TAG
I-TAG	A Backbone Service Instance TAG defined in the IEEE 802.1ah Standard [802.1ah]
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 IEEE 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. 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 both finished and ongoing work taking place in the IEEE 802.1 Working Group, the ITU-T, 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 Ethernet header fields: Virtual Local Area Network (VLAN) Identifiers (VIDs) 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.

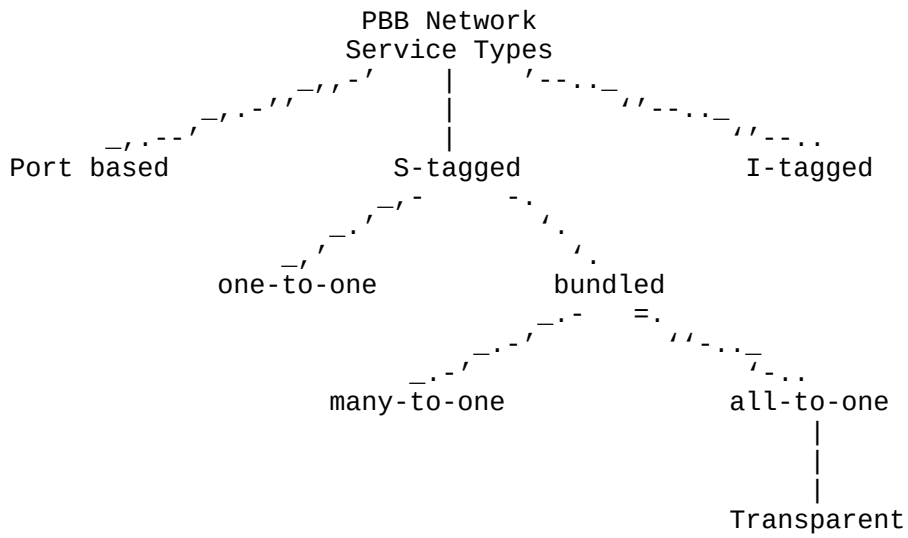

Figure 1: Ethernet Switching Service Types

The switching types are defined in Clause 25 of [802.1ah]. 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 (with the exception of the newly defined I-tagged service type).

[802.1ah] defines a new I-tagged service type but does not specifically define the Ethernet services being defined in the context of [MEF.6] and [G.8011], which are also illustrated in Figure 1.

To summarize the definitions:

o Port based

  This is a frame-based service that supports specific frame types; no Service VLAN tagging or MAC-address-based switching.

o S-tagged

  There are multiple 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 includes:

* 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 fields. 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 switching types determine which of the Ethernet header fields are used in the forwarding/switching function, e.g., VID only or VID and DMACs. The switching type may also require the use of additional Ethernet headers or fields. Services defined for UNIs tend to use the headers for requesting service (service delimiter) and are relevant between the customer site and network edge.

In most bridging cases, the header fields cannot be changed, but some translations of VID field values are permitted, typically at the network edges.

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

Ethernet OAM messages ([802.1ag] and [Y.1731]) rely on data-plane forwarding for both directions. Determining a broken path or misdirected packet in this case relies on OAM following the Eth-LSP. These OAM message identifiers are dependent on the data plane, so they work equally well for provisioned or GMPLS-controlled paths.

Ethernet OAM currently consists of:

Defined in both [802.1ag] and [Y.1731]:
- CCM/RDI: Continuity Check Message / Remote Defect Indication
- LBM/LBR: Loopback Message/Reply
- LTM/LTR: Link Trace Message/Reply
- VSM/VSR: Vendor-Specific Message/Reply

Additionally defined in [Y.1731]:
- AIS: Alarm Indication Signal
- LCK: Locked Signal
- TST: Test
- LMM/LMR: Loss Measurement Message/Reply
- DM: Delay Measurement
- DMM/DMR: Delay Measurement Message/Reply
- EXM/EXR: Experimental Message/Reply
- APS, MCC: Automatic Protection Switching, Maintenance Communication Channel

These functions are supported across all the standardized Eth-LSP formats.

2.3. Ethernet Switching Characteristics

Ethernet is similar to MPLS as it encapsulates different packet and frame types for data transmission. In Ethernet, the encapsulated data is referred to as MAC client data. The encapsulation is an Ethernet MAC frame with a header, a source address, a destination

address, and an optional VLAN identifier, type, and length on the front of the MAC client data with optional padding and a Frame Check Sequence at the end of the frame.

The type of MAC client data is typically identified by an "Ethertype" value. This is an explicit type indication, but Ethernet also supports an implicit type indication.

Ethernet bridging switches based on a frame's destination MAC address and VLAN. The VLAN identifies a virtual active set of bridges and LANs. The address is assumed to be unique and invariant within the VLAN. MAC addresses are often globally unique, but this is not necessary for bridging.

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 GMPLS architecture, per [RFC3945], allowed for control of Ethernet bridges and other Layer 2 technologies using the Layer-2 Switch Capable (L2SC) switching type. But, the control of Ethernet switching was not explicitly defined in [RFC3471], [RFC4202], or any other subsequent GMPLS reference document.

The GMPLS 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;
- forwarding priority level;
- connection preemption characteristics;
- protection/resiliency capability;
- routing policy, such as an explicit route;
- bidirectional service;

- end-to-end and segment protection;
- hierarchy

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

In applying GMPLS to "transport" Ethernet, GMPLS will need to be extended to work with the Ethernet data plane and switching functions. The definition of GMPLS support for Ethernet is multifaceted 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 characteristics of the "transport" Ethernet data plane are not modified in order to apply GMPLS control. For example, consider the IEEE 802.1Q [802.1Q] data plane: 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 made based on the DMAC. When forwarding using a 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. A special case is when a VID is defined for only two ports on one bridge, effectively resulting in a P2P forwarding constraint. In this case, all frames that are tagged with that VID and received over one of these ports are forwarded over the other port without address learning.

[802.1Qay] allows for turning off learning and hence the broadcast mechanism that provides means to create explicitly routed Ethernet connections.

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 a specific Ethernet service. Depending on the requirements of a service, it may be necessary to define multiple GMPLS protocol extensions and procedures. It is expected that all such extensions will be consistent with this document.

It is expected that a key requirement for service-specific documents will be to describe label formats and encodings. It may also be necessary to provide a mechanism to identify the required Ethernet service type in signaling and a way to advertise the capabilities of Ethernet switches in the routing protocols. These mechanisms must make it possible to distinguish between requests for different paradigms including new, future, and existing paradigms.

The Switching Type and Interface Switching Capability Descriptor share a common set of values and are defined in [RFC3945], [RFC3471], and [RFC4202] as indicators of the type of switching that should ([RFC3471]) and can ([RFC4202]) be performed on a particular link for an LSP. The L2SC switching type may already be used by implementations performing Layer 2 Switching including Ethernet. As such, and to allow the continued use of that switching type and those implementations, and to distinguish the different Ethernet switching paradigms, a new switching type needs to be defined for each new Ethernet switching paradigm that is supported.

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. Support of Eth-LSPs is expected to strictly comply to the GMPLS protocol suite addressing as specified in [RFC3471], [RFC3473], and related documents.

4.1. GMPLS Routing

GMPLS routing as defined in [RFC4202] uses IP routing protocols with 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 based on resource information obtained from LMP or from configured information. 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 (ISC) Descriptor. As mentioned in Section 3, the definition of one or more new ISCs to support Eth-LSPs is expected. Again, the L2SC ISCs will not be used to represent interfaces capable of supporting Eth-LSPs defined by this document and subsequent documents in support of the transport Ethernet switching paradigms. In addition, ISC-specific TE information may be defined as needed to support the requirements of a specific Ethernet Switching Service Type.

GMPLS routing is an optional functionality 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 connectivity network of sufficient capacity to handle the information exchange of the GMPLS routing and signaling protocols is necessary.

One way to implement this is with an IP-routed network supported by 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, or for forwarding user traffic over this IP network.

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

5. GMPLS Signaling

GMPLS signaling ([RFC3471] and [RFC3473]) is well suited to the control of Eth-LSPs and Ethernet switches. Signaling provides the ability to dynamically establish a path from an ingress node to an 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.

GMPLS signaling supports the establishment and control of bidirectional and unidirectional data paths. Ethernet is bidirectional by nature and CFM has been built to leverage this. Prior to CFM, the emulation of a physical wire and the learning requirements also mandated bidirectional connections. Given this,

Eth-LSPs need to be bidirectional congruent. 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. [RFC5467], an Experimental document, provides procedures if asymmetric bandwidth bidirectional LSPs are required.

6. Link Management

Link discovery has been specified for links interconnecting IEEE 802.1 bridges in [802.1AB]. 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. However, the 802.1AB capability is an optional feature, so it is not necessarily operating before a link is operational, and it primarily supports the management plane.

In the GMPLS context, LMP [RFC4204] has been defined to support GMPLS control-plane link management and discovery features. LMP also supports the automated creation of unnumbered interfaces for the control plane. If LMP is not used, there is an additional configuration requirement for GMPLS link identifiers. For large-scale implementations, LMP is beneficial. LMP also has optional fault management capabilities, primarily for opaque and transparent network technology. With IEEE's newer CFM [802.1ag] and ITU-T's capabilities [Y.1731], this optional capability may not be needed. It is the goal of the GMPLS Ethernet architecture to allow the selection of the best tool set for the user needs. The full functionality of Ethernet CFM should be supported when using a GMPLS control plane.

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 expected that LMP would be used for control-plane functions of link property correlation, but that Ethernet mechanisms for OAM such as CFM, link trace, etc., would be used for data-plane fault management and fault trace.

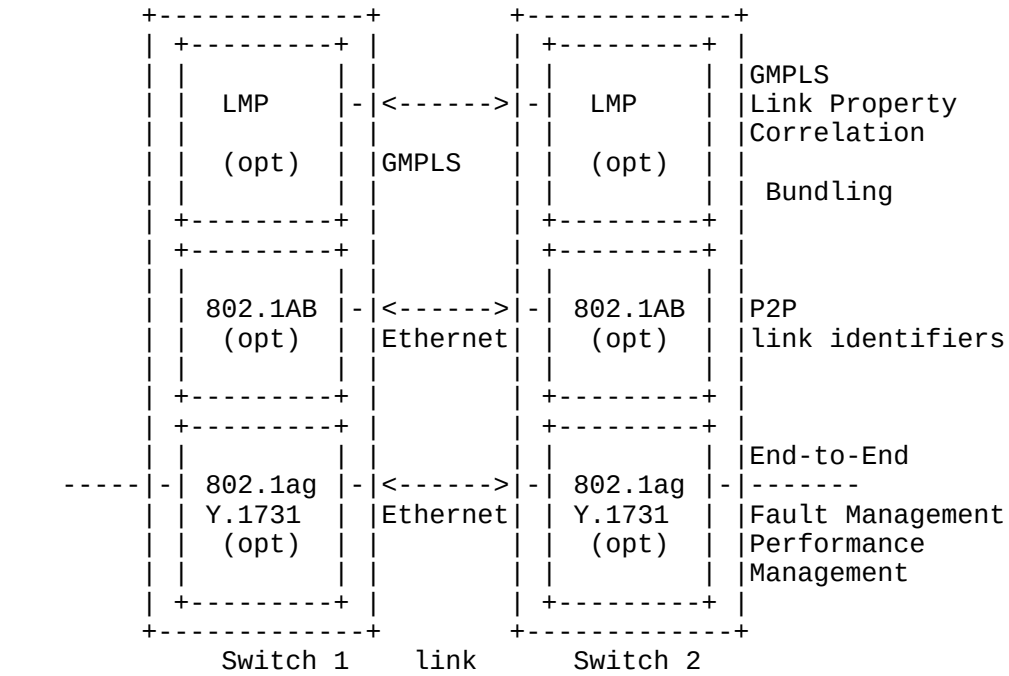

Figure 2: Logical Link Management Options

7. Path Computation and Selection

GMPLS does not identify a specific method for selecting paths or supporting path computation. GMPLS allows for a wide range of possibilities to be supported, from very simple path computation to very elaborate path coordination where a large number of coordinated paths are required. Path computation can take the form of paths being computed in a fully distributed fashion, on a management station with local computation for rerouting, or on 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 element and selection work; see [RFC4655].

8. Multiple VLANs

This document allows for the support of the signaling of Ethernet parameters across multiple VLANs supporting both contiguous Eth-LSP and Hierarchical Ethernet LSPs. The intention is to reuse GMPLS hierarchy for the support of peer-to-peer models, UNIs, and NNIs.

9. Security Considerations

A GMPLS-controlled "transport" Ethernet system should assume that users and devices attached to UNIs may behave maliciously, negligently, or incorrectly. Intra-provider control traffic is trusted to not 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 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, and 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 attacks may be required on UNI ports.

For a more comprehensive discussion on GMPLS security please see the MPLS and GMPLS Security Framework [SECURITY]. Cryptography can be used to protect against many attacks described in [SECURITY]. One option for protecting "transport" Ethernet is the use of 802.1AE Media Access Control Security [802.1AE], which provides encryption and authentication. It is expected that solution documents will include a full analysis of the security issues that any protocol extensions introduce.

10. References

10.1. Normative References

[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.

[RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.

10.2. Informative References

[802.1AB] "IEEE Standard for Local and Metropolitan Area Networks, Station and Media Access Control Connectivity Discovery", IEEE 802.1AB, 2009.

[802.1AE] "IEEE Standard for Local and metropolitan area networks Media Access Control (MAC) Security", IEEE 802.1AE-2006, August 2006.

[802.1ag] "IEEE Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks - Amendment 5: Connectivity Fault Management", IEEE 802.1ag, 2007.

[802.1ah] "IEEE Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks - Amendment 6: Provider Backbone Bridges", IEEE Std 802.1ah-2008, August 2008.

[802.1Q] "IEEE standard for Virtual Bridged Local Area Networks", IEEE 802.1Q-2005, May 2006.

[802.1Qay] "IEEE Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks - Amendment 10: Provider Backbone Bridge Traffic Engineering", IEEE Std 802.1Qay-2009, August 2009.

[ETH-TSPEC] Papadimitriou, D., "Ethernet Traffic Parameters", Work in Progress, January 2010.

[G.8011] ITU-T Recommendation G.8011, "Ethernet over Transport - Ethernet services framework", January 2009.

[G.8011.1] ITU-T Recommendation G.8011.1/Y.1307.1, "Ethernet private line service", January 2009.

[G.8011.2] ITU-T Recommendation G.8011.2/Y.1307.2, "Ethernet virtual private line service", January 2009.

[G.8031] ITU-T Recommendation G.8031, "Ethernet linear protection switching", November 2009.

[MEF.6] The Metro Ethernet Forum MEF 6, "Ethernet Services Definitions - Phase I", 2004.

[RFC4204] Lang, J., Ed., "Link Management Protocol (LMP)", RFC 4204, October 2005.

[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.

[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.

[RFC5467] Berger, L., Takacs, A., Caviglia, D., Fedyk, D., and J. Meuric, "GMPLS Asymmetric Bandwidth Bidirectional Label Switched Paths (LSPs)", RFC 5467, March 2009.

[SECURITY] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work in Progress, October 2009.

[Y.1731] ITU-T Recommendation Y.1731, "OAM Functions and Mechanisms for Ethernet based Networks", February 2008.

11. Acknowledgments

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

George Swallow contributed significantly to this document.

Authors' Addresses

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

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

Loa Andersson
Ericsson
Phone: +46 10 717 52 13
EMail: loa.andersson@ericsson.com