A Layered Network Architecture and Implementation for Ethernet Services

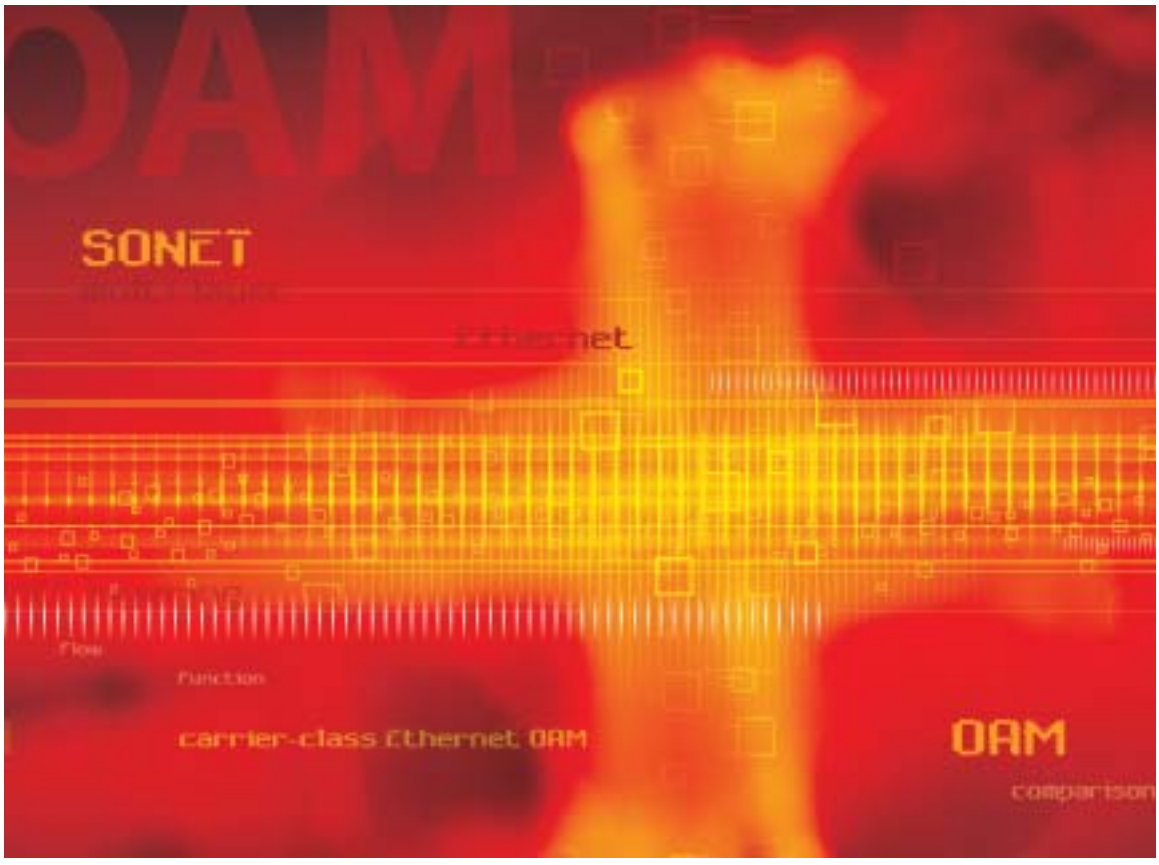

Introduction

Ethernet services are growing in popularity as end user services and as means to connect data sites within service provider networks. Customers like Ethernet services for their low-cost, high-bandwidth connectivity and ease of use. Carriers want to deploy Ethernet services to generate new revenues and to secure the business of enterprise customers. For these reasons, Ethernet services are expected to enjoy significant growth over the next several years; some market research suggests that, in North America, Ethernet service revenues may reach $4 billion by 2006.

Ethernet services are unlike previous data services: They use different physical interfaces (based on IEEE 802.3 Ethernet [1], as opposed to PDH or SONET) and they can support the broadcast functions inherent to Ethernet LANs. These new service attributes mandate new network architectures to support Ethernet services. Some small, early deployments of Ethernet feature flat networks of Ethernet CPE and carrier switches connected with fiber—networks that blend switching and transport. As services scale to support thousands of customers and billions of dollars in services revenue, the role of service-transparent transport (or, what the data community calls tunneling) becomes more important. A distinct transport function allows service providers to optimize the costs of transmission versus service switching and the operation and management of their networks.

Ethernet services operate at Layer 2. They may be tunneled using techniques at Layers 1, 2 or 3. Figure 1 illustrates the protocols that may provide transport for Ethernet services and their corresponding control planes.

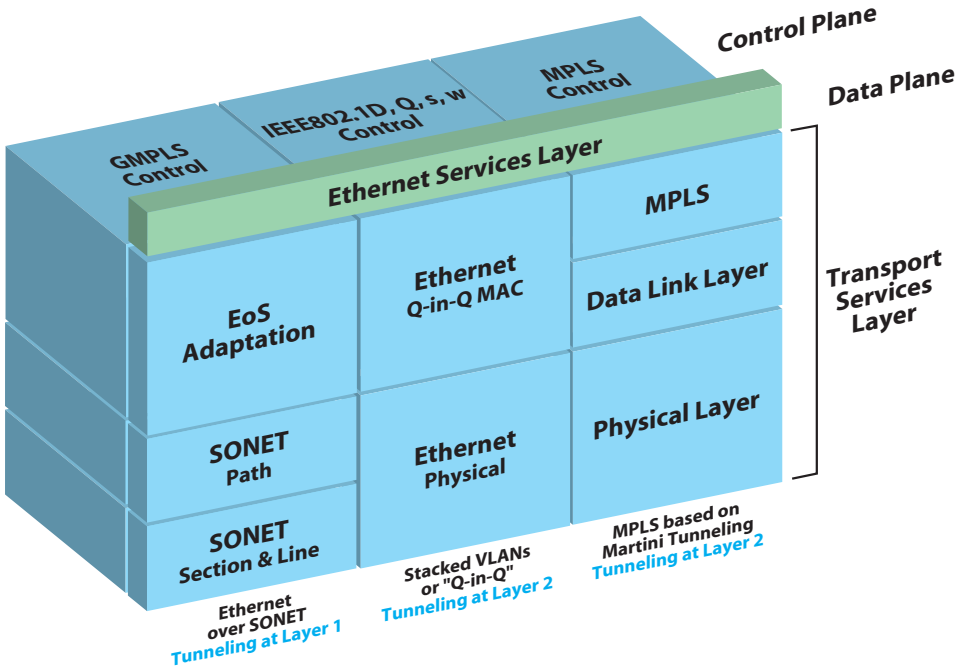

Figure 1: Ethernet Transport Protocols

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom

For your convenience, a list of acronyms can be found at the end of this document.

Investigating each of these transport alternatives for Ethernet services, this paper addresses issues such as survivability, standards compliance, interoperability, management, control and scalability. We will also discuss how each of these alternatives meets the evolving needs of service providers—from growing a nascent business under severe capital constraints to operating and managing a large-scale service.

This paper concludes with a proposal for a new kind of NE. This NE provides survivable, interoperable, manageable, and standards-compliant transport for Ethernet services, and does so in a way that integrates with the enormous installed base of SONET equipment and edge routers.

Reference Model

Figure 2 illustrates a reference model for a portion of a network that supports Ethernet services. Industry organizations such as the MEF [2] are developing detailed architecture models for Ethernet networks. The model presented here is an abstraction of those models intended to highlight the role of transport tunneling to support Ethernet services.

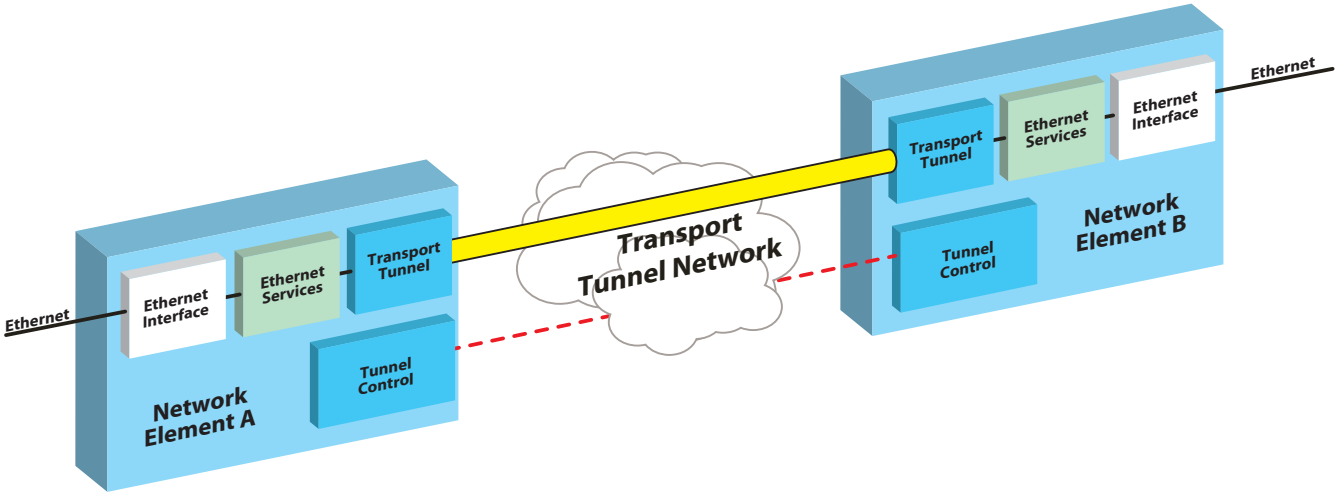

Figure 2: Reference Model

The model in Figure 2 shows two NEs that provide Ethernet services, along with a transport tunnel that connects these two NEs. As Ethernet frames flow from left to right in Figure 2, they first encounter the Ethernet interface functions of NE A. These functions provide IEEE 802.3-compliant physical and MAC layer functions.

Next, the Ethernet Services block of NE A provides the functions necessary to support the service associated with the Ethernet frame. This block corresponds the MEF's Ethernet Services Layer [2] and also to the Ethernet Services Layer in Figure 1. This functional block uses information in the customer's Ethernet frame, as well as provisioned information, to determine the functions necessary to support the Ethernet Line (E-Line) or Ethernet LAN (E-LAN) service [3] to which the customer has subscribed, and to perform these functions. These parameters could include bandwidth profile enforcement, Ethernet control protocol processing (if the Ethernet frame is a control PDU), QoS handling and determination of the next hop(s). In this reference model, the next hop is NE B.

The Ethernet frame proceeds to the transport tunnel functions associated with the tunnel that connects NE A with NE B. In this paper transport tunnel represents the logical link that connects adjacent Ethernet services entities. A transport tunnel operates a layer below the Ethernet Services Layer, providing services to the Ethernet Services Layer. The transport tunnel also has data, control and management planes that can operate independently from the data, control and management planes of the Ethernet Services Layer.

The transport tunnel functions reside at the Transport Services Layer (see [2] and Figure 1) and include processing of the protocols associated with the transport tunnel (which may include adding tunnel-specific protocol information) and receiving and transmitting on the physical transmission medium. Also shown in Figure 2 is the tunnel control function which supplies control plane functionality such as the signaling necessary to set up, supervise, and release connections and associated flows [2].[1]

The original Ethernet frame, possibly with tunnel control information added to it, then traverses the Transport Tunnel Network. This network operates only at the Transport Services Layer and performs no Ethernet Services Layer functions. The network may be modeled as a server to its client, the Ethernet Services Layer function [2]. Finally, the Ethernet frame arrives at NE B, where it traverses the same functional blocks as it did in NE A, except this time in the reverse order.

1 Figure 2 does not show management plane functions. Refer to a companion paper [4] for more details.

Transport Tunneling Techniques

Layer 1 Tunnels

In this instance, Layer 1 tunnels use SONET as the tunneling technique.[2] For this reason, they require an adaptation of Ethernet into SONET. Figure 3 illustrates the network and protocol models for Layer 1 tunnels.

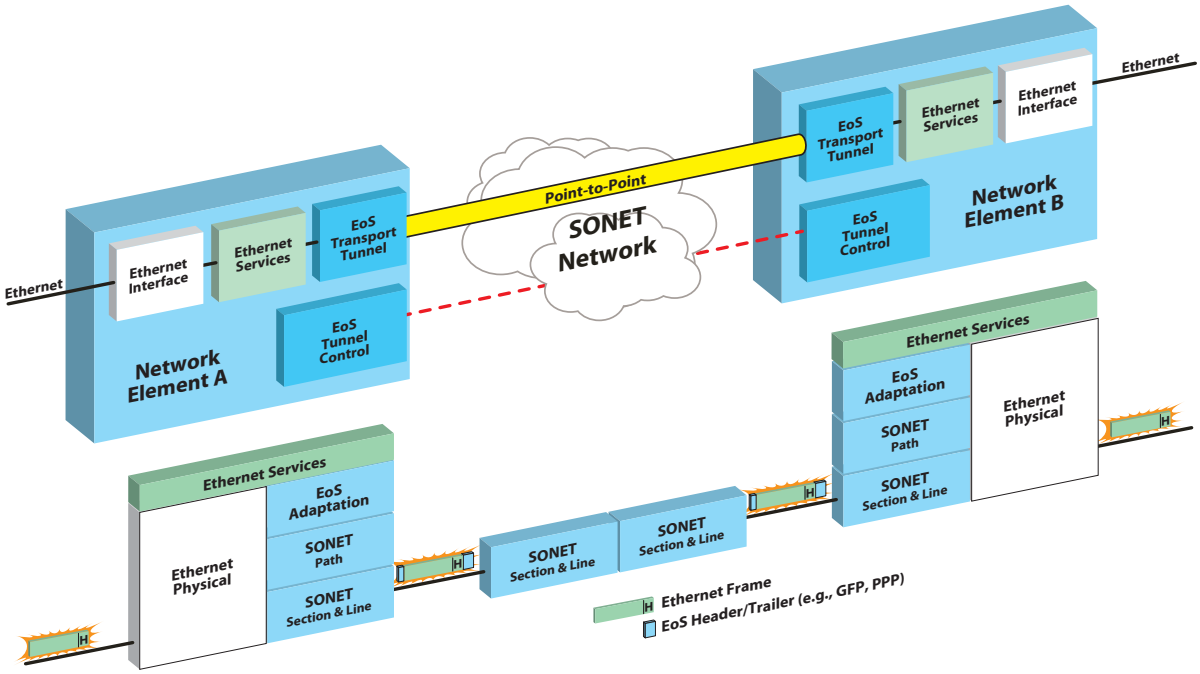

Figure 3: Tunneling of Ethernet at Layer 1

In this model, a SONET path or a virtual concatenation of SONET paths provides the point-to-point tunnel. Traditional SONET paths include STS-1 (51.84 Mbps), STS-nc (e.g., 622.08 Mbps for STS-12c) and VT1.5 (1.728 Mbps) paths. Virtual concatenation [6] combines a number of like paths (e.g., five STS-3c paths, virtually concatenated into an STS-3c-5v) to present a single payload to the EoS Adaptation Layer—a variation of inverse multiplexing of SONET paths into a single transport tunnel. Virtual concatenation provides additional bandwidth granularity for tunnels (i.e., at integer multiples of traditional SONET rates) in a manner that is transparent to the SONET network, since virtual concatenation is visible only at the SONET PTE, and the SONET network operates at the SONET Section and Line Layers (Refer to Figure 3).

SONET interfaces support physical layer channelization—the ability to multiplex STS or VT paths onto a single physical interface. If a Layer 1 tunnel is a SONET path or virtual concatenation of SONET paths, then channelized SONET interfaces (e.g., an OC-12 interface with twelve STS-1 paths) generally support the ability to carry multiple Layer 1 tunnels. Concatenated interfaces (e.g., an OC-12 interface with a single STS-12c path) carry a single Layer 1 tunnel.

2 This comparison does not consider WDM technology, which could also be considered a Layer 1 tunneling technique.

SONET paths require an EoS adaptation to carry Ethernet frames, primarily to handle the framing of Ethernet within the SONET payload. Three standard framing methods exist: GFP [7], PPP [8,9] and X.86 [10]. All of these methods add header information and, in some cases, trailer information to each Ethernet frame, as Figure 3 illustrates. Other EoS adaptation functions include rate adaptation (i.e., matching the Ethernet interface rate with the rate of the SONET payload) and OAM adaptation (i.e., mapping between Ethernet OAM and SONET OAM).

Additional attributes of Layer 1 tunnels:

- **Survivability** – SONET offers a variety of protection mechanisms that provide restoration in fewer than 50 ms following detection of a failure. These mechanisms include linear configurations (e.g., 1+1 and 1:1) as well as UPSR and BLSR mechanisms. Most protection mechanisms in use today require the network to set aside half the bandwidth as protection bandwidth.

- **Standards and Interoperability** – Several established standards, most notably Telcordia® GR-253-CORE [11], define SONET. While service providers have deployed few multi-vendor rings, SONET interfaces are ubiquitous as high capacity meet-points between equipment from different vendors and different service providers. Virtual concatenation standards are also complete [6]. Standardization efforts for EoS vary in their degrees of maturity. Framing standards, although nascent, are complete. ITU-T SG 13 has only recently begun work on the many of the operational aspects of EoS adaptation. In 2001, the MEF began work on an EoS interoperability agreement to define a set of common options to allow standard EoS implementations to interoperate; that effort remains unfinished.

- **Management** – SONET technology offers a standard set of operations capabilities, including performance monitoring and fault surveillance. Most major service providers have deployed intricate OSSs that enable them to use these capabilities in large-scale networks. Most SONET systems use TL1 as the OSS management protocol. Adding Ethernet functions can complicate the management model, since legacy SONET OSSs typically do not recognize Ethernet switching capabilities and Ethernet NE management is usually defined in the context of SNMP. However, treating the Ethernet capabilities as a transport tunnel helps to mitigate some of these difficulties because, with relatively little effort, legacy transport OSSs can be upgraded to support these new point-to-point circuits at a considerable investment by equipment vendors.

- **Control** – Figure 3 illustrates the control plane functions for Layer 1 tunnels. In principle, service providers may use GMPLS for control of EoS-based Layer 1 tunnels. For a host of complicated reasons (many centered on compatibility with the existing operations infrastructures), most service providers use OSS-based control. Virtual concatenation presents an additional consideration for the control plane, since it requires some control of the grouping of SONET paths into a single Layer 1 tunnel. LCAS [12] defines a method to dynamically control the membership of a transport tunnel comprising virtually concatenated SONET paths. LCAS does not provide connection management for the constituent paths. At the tunnel endpoints, LCAS handles the addition and deletion of existing SONET paths to the tunnel.

- **Scalability** – Scaling connectivity in SONET networks can prove difficult for three reasons: (1) They are connection-oriented, (2) SONET connections are fixed-bandwidth, and (3) typically, SONET networks use no automated control plane. These factors limit the use of Layer 1 tunnels in the core, although they have great utility in access networks, where SONET is the predominant optical access technology, and scalability is less of a concern.

Layer 2 Tunnels

Ethernet is inherently a LAN technology, which makes it difficult to view Ethernet as a tunnel technology. Techniques such as stacked VLANs allow creation of an Ethernet network that operates a layer below the Ethernet Services Layer. This Ethernet network has its own user, control, and management planes, which operate independently from those at the Ethernet Services Layer—which is precisely the role of a tunnel network in a layered network architecture. Figure 4 illustrates the network and protocol models for Layer 2 tunnels.

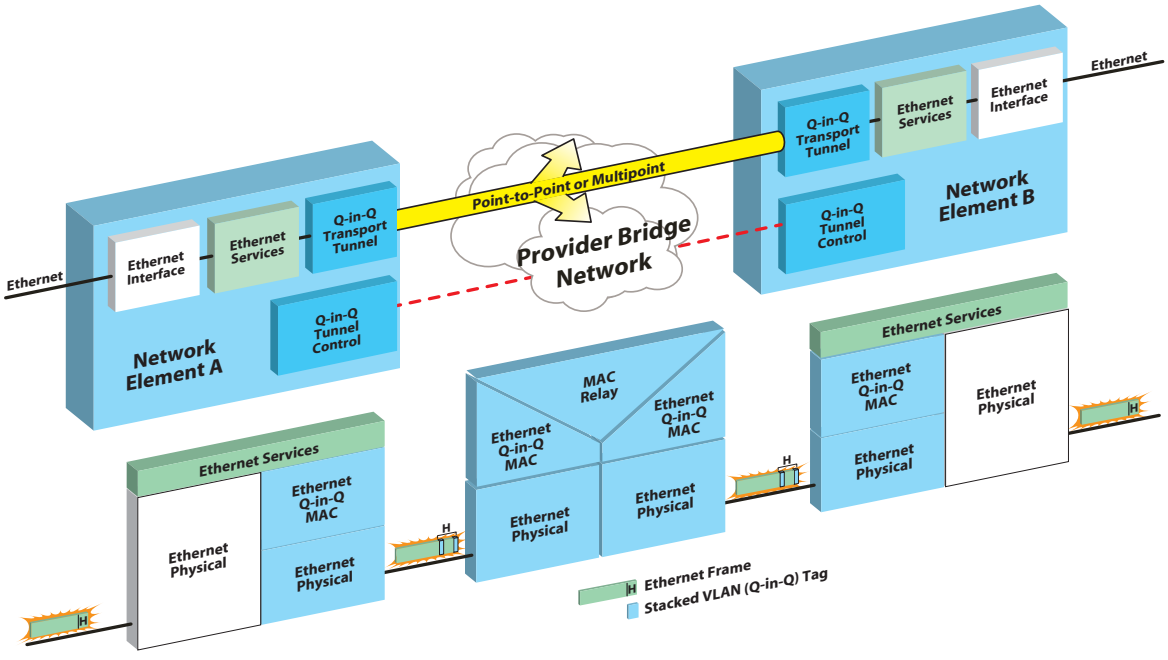

Figure 4: Tunneling of Ethernet at Layer 2

In this model, the Layer 2 tunnels comprise stacked VLANs, or what IEEE 802.1ad [13] terms a P-VLAN. Ethernet frames at NE A that, according to the determination made by the Ethernet services function, must reach the Ethernet services function at NE B are assigned to a P-VLAN that includes the corresponding Ethernet services function at NE B as a member. NE A adds to each of these frames an additional VLAN tag that identifies the P-VLAN (see Figure 4). The Provider Bridge network [13] uses this new outer VLAN tag, along with the original Ethernet DA, to transport each of these frames to the Ethernet Services function at NE B.

As the name suggests, a P-VLAN provides the functions of a LAN, including broadcast capabilities. A P-VLAN can be defined with multiple members—with multiple Ethernet services layer instances, possibly at multiple NEs. The P-VLAN may broadcast frames with broadcast Ethernet DAs to all members of the P-VLAN; moreover, frames with unknown DAs are also broadcast to all members of the P-VLAN, so that the address may be learned. This broadcast capability differentiates Layer 2 tunnels from Layer 1 tunnels, and can act as a useful tool in building networks to support multipoint Ethernet services.

Figure 4 shows the Layer 2 tunnel with an Ethernet Physical layer. One of the distinguishing features of these tunnels is that they operate at Layer 2 and may use any standard Layer 1 technology, including SONET.

Additional attributes of Layer 2 tunnels:

- **Survivability** – Layer 2 tunnels rely on Spanning Tree techniques such as RSTP [14] and MSTP [15] for network restoration. These methods place no upper bound on restoration time (unlike the 50 ms of SONET). Convergence times can range from a few hundred milliseconds to several seconds, depending on the network configuration.

- **Standards and Interoperability** – Ethernet bridging is based on the venerable IEEE 802.1D [16] and 802.3 standards; IEEE 802.1Q [17] defines VLAN bridging capabilities. In 2002, the IEEE 802.1 working group (project IEEE 802.1ad) began the effort to standardize provider bridges. The working group's incremental goal is to enable "service providers to use the architecture and protocols of 802.1Q" [13]—and will be documented in an Amendment to IEEE 802.1Q. Many Ethernet switch vendors support pre-standard implementations of provider bridging. User plane interoperability is straightforward, and centers on consistent interpretation of the P-VLAN tags. Control plane issues (e.g., how does a provider bridge network handle users' control frames?) pose greater interoperability challenges. IEEE 802.1ad will address these issues, although many of them require further study [13].

- **Management** – Ethernet technology also offers a rich set of management capabilities that focus mainly on the management of nodes and links. However, Ethernet's inherent broadcast nature and its lack of a path concept make end-to-end service management difficult. Building Layer 2 tunnels using P-VLANs helps mitigate these difficulties, especially if the tunnels are point-to-point (i.e., a P-VLAN with two members). Ethernet switches typically use SNMP.

- **Control** – Ethernet switches have a fairly rich control plane that features automated address learning, protocols such as RSTP and MSTP for topology management, and GVRP for configuration and management of VLAN membership. The IEEE 802.1ad standard should allow provider bridges to extend these concepts to P-VLAN Layer 2 tunnels. As stated previously, perhaps the most prominent open issue is the interaction between users' control protocols and the control plane of the provider bridge network.

- **Scalability** – The P-VLAN tag comprises 12 bits, allowing a provider bridge network to support up to 4,094 P-VLANs (two values are reserved). Provider bridge networks also run into two scaling issues: MAC address table scaling, since each provider bridge must eventually learn every MAC address behind every P-VLAN that it supports; and Spanning Tree scaling. There is general industry consensus that Layer 2 networks afford limited scalability. The point at which scalability becomes a concern is a topic for lively industry debate.

Layer 3 Tunnels

While it may appear counterintuitive to use a Layer 3 technology to tunnel a Layer 2 service, such tunnels are in use today [18]. These tunnels use MPLS as the fundamental transport technology. In IETF terminology, these point-to-point Layer 3 tunnels are "VCs or 'pseudowires' that make use of underlying PSN tunnels" [5]. These PSN tunnels should not be confused with the transport tunnels described in this paper. PSN tunnels connect PE devices, which correspond to the NEs in Figure 5. Layer 3 tunnels (or VCs or pseudowires) ride over PSN tunnels to link Ethernet Services Layer entities within those NEs. Each PE (or NE) device may support multiple Ethernet Services Layer entities; a PSN tunnel may therefore carry multiple VCs/pseudowires/Layer 3 tunnels.

Figure 5 illustrates the network and protocol models for Layer 3 tunnels. The pseudowire layer adds an inner MPLS label. The PSN tunnel layer adds its own protocol information. The PSN tunnel may employ one of several different technologies. Figure 5 shows an MPLS-based PSN tunnel layer, so in this example, the PSN tunnel protocol information comprises an outer MPLS label. Since MPLS operates at Layer 3 it may use a variety of Layer 2 and Layer 1 protocols, as Figure 5 illustrates.

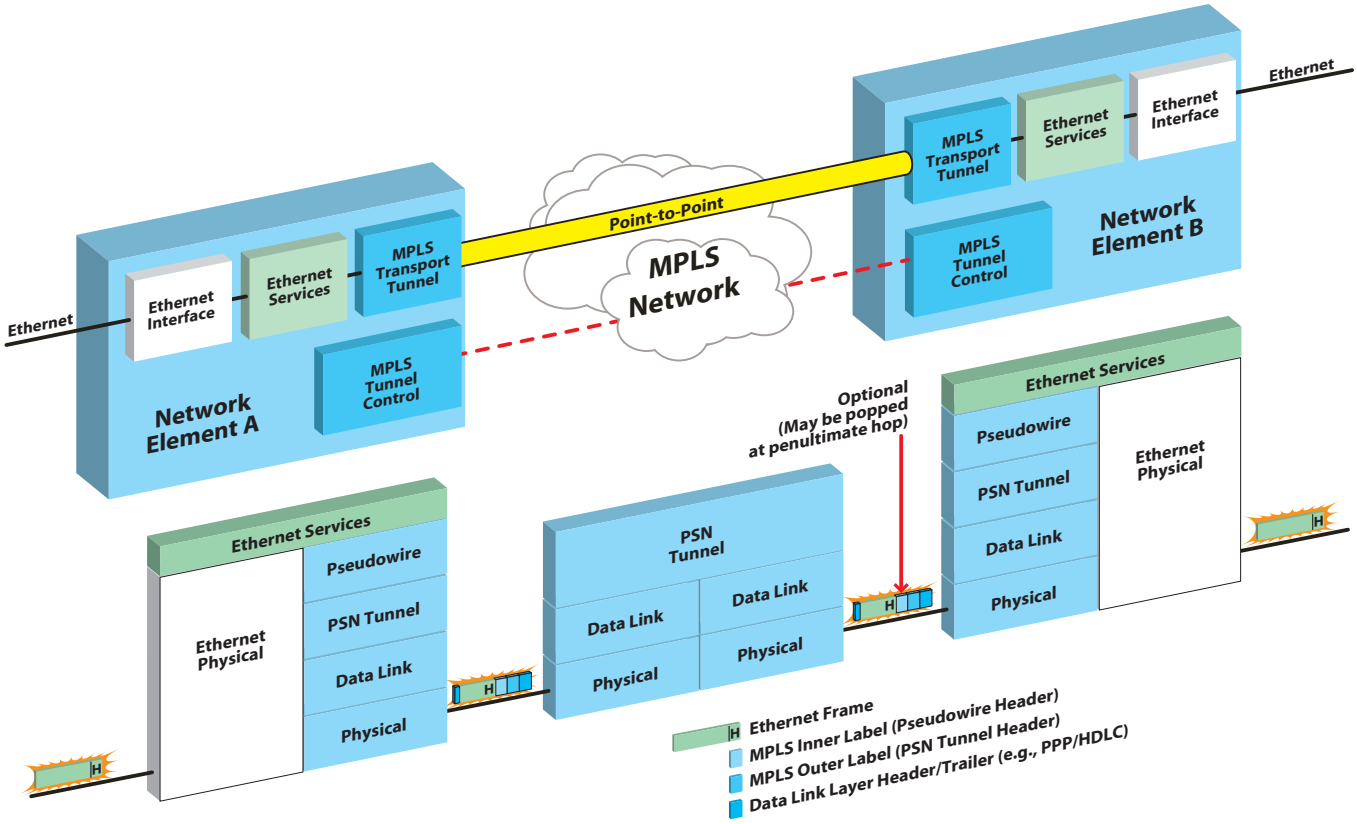

Figure 5: Tunneling of Ethernet at Layer 3

Additional attributes of Layer 3 tunnels:

- **Survivability** – MPLS supports an FRR capability, which enables the establishment of backup LSP tunnels for local repair of LSP tunnels. In the event of a failure, these backup tunnels allow redirection of traffic in tens of milliseconds [19]. A service provider typically would use the FRR capability to protect the PSN tunnels, since they define the logical network topology; Layer 3 tunnels (i.e., pseudowires) would ride on top of these protected PSN tunnels. While this work remains in draft form in the IETF [19], several vendors have begun implementing and testing interoperability of this feature [20].

- **Standards and Interoperability** – IETF can claim most of the standards work on MPLS. While technically not a standards body, the IETF has sanctioned a number of RFCs that define aspects of MPLS; from this perspective, the seminal work on MPLS may be considered mature. Other work specific to MPLS tunneling of point-to-point Ethernet services remains in draft form [5], although service providers have begun deploying services based on these drafts [18]. More recently, work in IETF has begun on issues such as tunneling for multipoint Ethernet service, or VPLS [21]. The ITU-T SG13 has also begun work on OAM for MPLS, and has produced several recommendations on this subject.

- **Management** – In principle, MPLS allows for end-to-end tunnel management, since it supports the notion of a path (e.g., a tunnel is an MPLS LSP). The ITU-T has begun standardizing OAM requirements and mechanisms for MPLS [22]. Moreover, the IETF is defining how to use underlying MPLS techniques (e.g., MPLS ping packets, MPLS signaling) to support Ethernet service-level OAM functions such as connectivity verification and topology discovery [23]. MPLS routers typically use SNMP.

- **Control** – The advanced IP-based control plane provides many of the principal benefits of MPLS. MPLS routers may use the LDP or RSVP to control LSPs. TE variants of these allow LSP control with additional constraints.

- **Scalability** – MPLS offers Internet-size scalability. MPLS allows for hierarchical aggregation—or LSPs within LSPs. While Figure 5 illustrates a two-level hierarchy, theoretically any number of LSPs may be stacked in this fashion. Moreover, the MPLS control plane uses IP addresses, and the data community understands well how to use this time-tested address structure to implement large networks (e.g., the Internet). For these reasons, MPLS often is viewed as a solution for scalability problems at Layer 1 and Layer 2.

Tunnel Comparison Summary

The table below summarizes the attributes of the Ethernet tunneling techniques at Layers 1, 2, and 3.[3]

Attribute	Layer 1 Tunnels	Layer 2 Tunnels	Layer 3 Tunnels
Protocols	• Ethernet over SONET • GFP, PPP, X.86 • Virtual Concatenation	• Ethernet, P-VLANs	• MPLS pseudowires • Underlying PSN tunnels may also be MPLS
Topology	• Point-to-point paths • Typically physical ring or star	• Point-to-point • Multipoint using Ethernet broadcast	• Point-to-point LSPs • Physical star, ring, mesh or combinations
Survivability	• SONET APS, UPSR,BLSR, all protection switches ≤ 50 ms	• RSTP, MSTP • Restoration time varies with network configuration	• MPLS FRR for PSN tunnels • 10's of ms restoration
Standards and Interoperability	• GR-253-CORE, T1.105, Interoperable SONET handoffs • EoS interop still immature	• IEEE 802.3, 802.1Q, 802.1ad (Provider Bridges) • Some pre-standard Q-in-Q interoperability	• Internet RFCs, MPLS drafts mature • Some point-to-point EoMPLS interop based on martini drafts
Management	• Robust OAM • Embedded OSSs • TL1 for OSS • Point-to-point EoS tunnels fit existing model	• Robust node, link OAM • Challenges: broadcast, no path concept SNMP for NE-OSS	• Path concept allows end-to-end OAM • Work in early stages in ITU-T, IETF • SNMP for NE-OSS
Control	• OSS (today), GMPLS (tomorrow) • LCAS for control of virtual concatenation	• RSTP, MSTP, GVRP • IEEE 802.1ad will likely extend to Provider Bridges	• Major benefit: RSVP, LDP TE extensions
Scalability	• Limited due to fixed bandwidth • Connection-oriented tunnels and no control plane (today)	• 4,094 per P-VLANs • MAC address and Spanning Tree scaling issues	• Scales to Internet-size • Hierarchical aggregation • IP addressing
Where to Use Ethernet Tunneling Technique	• Access: With SONET installed base or for circuit/Ethernet combination • IOF: As a physical layer for other tunnels	• Metro IOF: Multipoint capabilities provide efficiency; likely no scalability issues	• Ethernet network core due to unparalleled scalability

From this table, and from the preceding discussion, we make two important observations:

1. While the three technologies use very different protocols and work at different layers, each is a valid tunneling method (i.e., Transport Services Layer) for Ethernet services.
2. Each of the tunneling methods has unique attributes that allow it to function well in a particular part of the network. For example, Layer 1 tunnels are well suited for access networks where SONET is the predominant optical access technology. Layer 2 tunnels can prove useful in the middle of the network where their broadcast capabilities may enable efficient transport for multipoint services and where scalability is not an issue. Finally, Layer 3 tunnels are ideally suited for the core of Ethernet service networks, since they provide unparalleled scalability.

3 There is a possibility to combine tunneling methods (e.g., Ethernet over MPLS over SONET) to take advantages of the beneficial attributes of each method. The following section describes an implementation that uses tunneling methods in combination.

Ethernet Transport Tunnel Manager

Figure 6 depicts a metro network that supports Ethernet services. The network comprises several different kinds of NEs, each with a specific role in the network. These NEs include MSPPs, which provide access to Ethernet services over EoS (Layer 1 tunnels); Ethernet MCs, which provide native Ethernet access to Ethernet services (no tunnels); and provider bridges, which use Layer 2 tunnels (possibly multipoint) in the interoffice network. Several nodes within the network also support Ethernet Services Layer functions in addition to Ethernet transport tunnels.

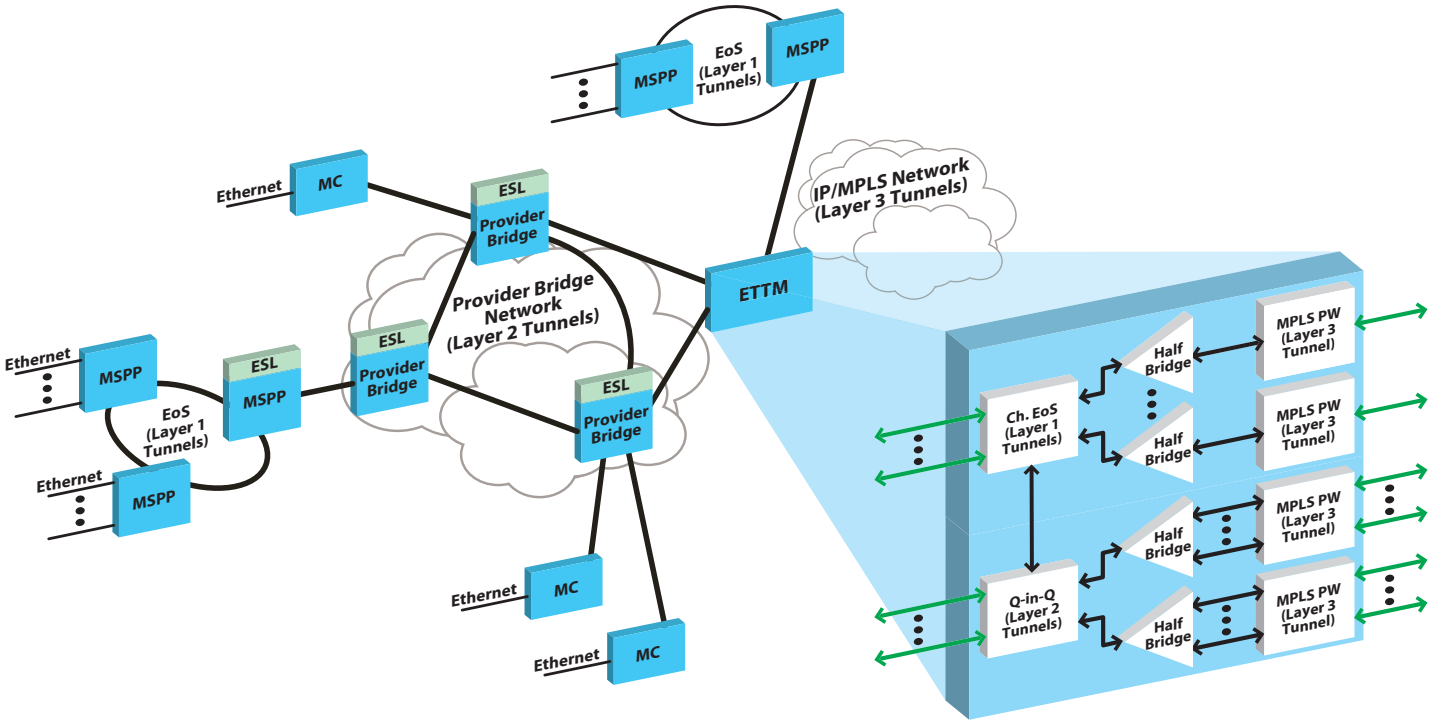

Figure 6: Ethernet Transport Tunnel Mapper

Figure 6 also shows a new kind of NE: An ETTM. This node extends Layer 1 and Layer 2 tunnels over Layer 3 tunnels. An ETTM is important as a distinct NE because it provides the transition from metro networks (access and IOF) to core networks; and it allows the relatively simple control planes at Layers 1 and 2 to operate over the extensive Layer 3 control plane. The ETTM peers with edge and core routers in the IP/MPLS network.

The ETTM must be able to support many of the Layer 1, 2, and 3 tunnel capabilities (e.g., Ethernet over channelized SONET, P-VLANs on Ethernet interfaces, MPLS pseudowires) described in this paper. In addition, the ETTM uses a unique half-bridge model to extend even multipoint Layer 2 tunnels over the IP/MPLS core network. The ETTM supports a half-bridge for each Layer 1 or Layer 2 tunnel; each half-bridge maps the Layer 1 or 2 tunnel into one or more Layer 3 tunnels. A half-bridge with multiple Layer 3 tunnels can transport a multipoint Layer 2 tunnel by observing three simple rules:

1. **Replicate across the Layer 3 tunnels** – The half-bridge replicates Ethernet frames with broadcast or unknown MAC addresses across all the Layer 3 tunnels.
2. **Learn from the Layer 3 tunnels** – The half-bridge learns MAC addresses from Ethernet frames received on the Layer 3 tunnels.
3. **Split horizon** – The half-bridge never forwards Ethernet frames between Layer 3 tunnels.

Figure 7 illustrates the half-bridge operation. As a degenerate case, a half-bridge with a single Layer 3 tunnel performs a simple mapping of a point-to-point Layer 1 or 2 tunnel into a Layer 3 tunnel.

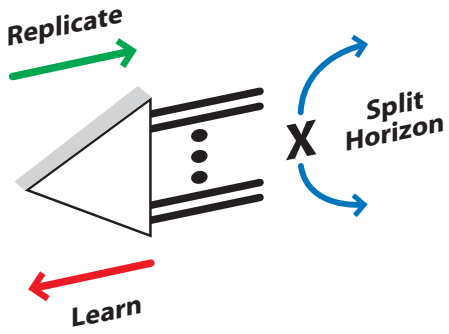

Figure 7: Half-Bridge Operation

Conclusion

Ethernet services continue to grow in popularity. Even though Ethernet is a Layer 2 service, it may be transported using tunnels at Layer 1 (EoS), Layer 2 (P-VLANs), or even Layer 3 (MPLS pseudowires). When comparing tunneling techniques, we draw two broad conclusions: First, each of these tunneling techniques represent a valid method for transporting Ethernet services—even approaches that support multipoint tunnels (e.g., P-VLANs) or those that operate at Layer 3 (MPLS). Second, each of these approaches offers unique benefits and limitations that suit them for particular applications in service providers' networks.

An ETTM extends Layer 1 and Layer 2 tunnels (which do not scale well) over highly-scalable Layer 3 tunnels. The ETTM features a unique half-bridge implementation that, by adhering to three simple rules, allows it to support the transport of both point-to-point Layer 1 and Layer 2 tunnels and multipoint Layer 3 tunnels.

References

[1] *IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements*, Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, IEEE 802.3-2002.

[2] *Metro Ethernet Network Architecture Framework, Part 1: Generic Framework*, Metro Ethernet Forum, Approved Draft Version 2.0, April 15, 2003.

[3] *Ethernet Services Definitions, Phase 1*, Metro Ethernet Forum, Straw Ballot Draft 3.8, April 2, 2003.

[4] O'Connor, D., *Comparison of Ethernet and SONET OAM - Status Report on Carrier Class Ethernet OAM*, NFOEC 2003.

[5] *Transport of Layer 2 Frames over MPLS*, Martini et al., draft-martini-l2circuit-trans-mpls-11.txt, April 2003.

[6] Recommendation G.707, *Network Node Interface for the Synchronous Digital Hierarchy*, ITU-T, April 2003 (includes various revisions, corrigenda, and amendments).

[7] Recommendation G.7041/Y.1301, *Generic Framing Procedure*, ITU-T, March 2003 (includes Amendments 1 and 2 and Corrigendum 1).

[8] RFC 2615, *PPP over SONET/SDH*, A. Malis and W. Simpson, IETF, June 1999.

[9] RFC 1662, *PPP in HDLC-like Framing*, W. Simpson, IETF, July 1994.

[10] Recommendation X.86, *Ethernet over LAPS*, ITU-T, April 2002 (includes Amendment 1).

[11] *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria*, GR-253-CORE, Telcordia, Issue 3, September 2000.

[12] G.7042/Y.1305, *Link Capacity Adjustment Scheme (LCAS) for Virtual Concatenated Signals*, ITU-T, March 2003 (includes Amendment 1 and Corrigendum 2).

[13] *Draft IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks—Amendment 4: Provider Bridges*, IEEE P802.1ad/D1, May 14, 2003.

[14] *IEEE Standard for Local and Metropolitan Area Networks—Common Specifications, Part 3: Media Access Control (MAC) Bridges—Amendment 2: Rapid Reconfiguration*, IEEE Std 802.1w-2001.

[15] *Draft IEEE Standard for Local and Metropolitan Area Networks—Amendment 3 to 802.1Q Virtual Bridged Local Area Networks: Multiple Spanning Trees*, IEEE P802.1s/D15, October 9, 2002.

[16] *IEEE Standard for Local and Metropolitan Area Networks—Common Specifications, Part 3: Media Access Control (MAC) Bridges—Part 3: Media Access Control (MAC) Bridges*, IEEE Std 802.1D-1998.

[17] *IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks*, IEEE Std 802.1Q-1998.

[18] *Interview with Luca Martini, Level 3*, www.lightreading.com, May 15, 2003.

[19] *Fast Reroute Extensions to RSVP for LSP Tunnels*, Swallow et al., draft-ietf-mpls-rsvp-lsp-fastreroute-02.txt, February 2003.

[20] *MPLS Vendors Demo Fast Reroute*, www.lightreading.com, October 31, 2002.

[21] *Virtual Private LAN Services over MPLS*, Lasserre et al., draft-lasserre-vkompella-ppvpn-vpls-04.txt, March 2003.

[22] Recommendation Y.1711, *OAM Mechanism for MPLS networks*, ITU-T, November 2002.

[23] *Testing Hierarchical Virtual Private LAN Services*, Stokes et al., draft-stokes-vkompella-ppvpn-hvpls-oam-01.txt, December 2002.

Acronym	Descriptor
BLSR	Bidirectional Line Switched Ring
CPE	Customer Premises Equipment
DA	Destination Address
EoMPLS	Ethernet over MultiProtocol Label Switching
EoS	Ethernet over SONET
ESL	Ethernet Service Layer
ETTM	Ethernet Transport Tunnel Manager
FRR	Fast Reroute
GARP	Generic Attribute Registration Protocol
GFP	Generic Framing Protocol
GMPLS	Generalized MultiProtocol Label Switching
GVRP	GARP VLAN Registration Protocol
HDLC	High level Data Link Control
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IOF	Interoffice Facilities
IP	Internet Protocol
ITU	International Telecommunication Union
LAN	Local Area Network
LCAS	Link Capacity Adjustment Scheme
LDP	Label Distribution Protocol
LSP	Label Switched Path
MAC	Media Access Control
MC	Media Converter
MEF	Metro Ethernet Forum
MPLS	MultiProtocol Label Switching
MSPP	MultiService Provisioning Platform
MSTP	Multiple Spanning Tree Protocol

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom

Acronym	Descriptor
NE	Network Element
OAM	Operations, Administration and Maintenance
OSS	Operational Support System
PDH	Plesiochronous Digital Hierarchy
PDU	Packet Data Unit
PE	Provider Edge
PPP	Point-to-Point Protocol
PSN	Packet Switched Network
PTE	Path Terminating Element
P-VLAN	Provider–Virtual Local Area Network
PW	Pseudowire
QoS	Quality of Service
RFC	Request For Comment
RSTP	Rapid Spanning Tree Protocol
RSVP	Resource reSerVation Protocol
SNMP	Simple Network Management Protocol
SONET	Synchronous Optical Network
TE	Traffic Engineering
TL1	Transaction Language 1
UPSR	Unidirectional Path Switched Ring
VC	Virtual Circuit
VLAN	Virtual Local Area Network
VPLS	Virtual Private LAN Service
WDM	Wavelength Division Multiplexing

FUJITSU NETWORK COMMUNICATIONS INC.
2801 Telecom Parkway, Richardson, Texas 75082-3515
Telephone: (972) 690-6000
(800) 777-FAST (U.S.)
us.fujitsu.com/telecom