

# Virtual Private LAN Services: Extending Ethernet over the WAN

The market for data networking services is constantly evolving. Demand for bandwidth grows exponentially as end users continue to look for more sophisticated features and services. In the midst of this constant change, Ethernet has persevered and grown for several decades. Since the 1980s, Ethernet has evolved from a Local Area Network (LAN) protocol running at 10 Mbps over a shared coaxial cable to a ubiquitous networking technology running at speeds up to 10 Gbps over all manner of copper wire and optical fiber in networks extending worldwide. The continued acceptance and familiarity of Ethernet has led service providers to offer a variety of carrier Ethernet services to extend these networks beyond the LAN and into regional, national, and now international Wide Area Network (WAN) connections.

Carrier Ethernet has traditionally been offered as one of three service types:

- **Ethernet Access to IP Network:** Access from an enterprise customer site to a Layer 3 network such as an IP backbone or Layer 3 MPLS VPN. These services are frequently offered with Ethernet over SONET/SDH encapsulation. Both X.86 and Generic Framing Procedure (GFP) provide framing methods for encapsulating and mapping Ethernet frames into SONET/SDH infrastructure.
- **Point-to-Point Ethernet Private Line:** Services such as Ethernet Virtual Private Line (EVPL) and Virtual Private Wire Services (VPWS) are used to connect two customer locations with a point-to-point Ethernet link.
- **Transparent LAN Service (TLS):** A virtual Ethernet switch or switches in the service provider network that provides any-to-any connectivity among a group of remote enterprise locations. TLS has been most commonly deployed as Ethernet flat bridged networks using 802.1q VLAN tags to provide service instance separation.

Virtual Private LAN Services (VPLS) has now emerged as a more scalable, WAN-focused, multipoint Ethernet service type to connect multiple remote enterprise locations. VPLS provides a new technical solution for deploying TLS over a Multiprotocol Label Switching (MPLS) core network as described in this white paper.

## Virtual Private LAN Services

VPLS, as outlined in IETF RFC 4762 and approved in January of 2007, is the latest and most advanced IETF standard supporting carrier Ethernet service. VPLS provides a sophisticated architecture whereby multisite switched Ethernet service is offered over a converged MPLS backbone. This new hierarchical service offering allows carriers to design Ethernet services utilizing a flexible MPLS core, and expand this offer with a variety of new features and capabilities not traditionally available in TLS offers.




Figure 1. Three types of carrier Ethernet service

This white paper includes an overview of the technical building blocks of VPLS, including a review of the MPLS features that help enable the creation of secure customer networks within a shared packet infrastructure. The concept of pseudowires (PW) is covered with a brief description of Ethernet private line service provisioning over an MPLS core. This concept is expanded to a full mesh of PWs, which forms the heart of the any-to-any connectivity of VPLS. The paper concludes with a review of several advanced features, including QoS, Hierarchical VPLS (H-VPLS) and various interworking options.

## The Technical Foundation of VPLS

VPLS links together multiple sites in a single bridged Ethernet domain running over an MPLS network. This single Ethernet domain is created by building point-to-point virtual circuits or PWs, across MPLS Label Switched Path tunnels (LSP). PWs can be deployed either as point-to-point links, such as Virtual Private Wire Service (VPWS), or in a full mesh as in VPLS. The following section reviews the key technologies behind VPLS.

## Multiprotocol Label Switching

MPLS has evolved over the past decade as the carrier's technology of choice for supporting multiservice and IP networks. MPLS was originally designed to solve the traffic-engineering problems created by the explosive growth of the Internet traffic in the early 1990s. Rising volumes of multimedia, voice and data traffic put an enormous strain on the Asynchronous Transfer Mode (ATM) switches and software-based routers that comprised the Internet at that time. MPLS was developed as a faster, more efficient alternative to traditional IP routing. However, as router processing power increased and wire-speed hardware-based forwarding emerged as the industry norm, other benefits of MPLS, such as multiprotocol support, QoS, traffic engineering features and segmentation of private networks over a shared infrastructure, eventually became equally compelling features driving MPLS adoption across the industry.

Several characteristics differentiate MPLS from simple IP routing to allow the creation of Ethernet service offerings:

### Label Stacking

In MPLS implementations, physical or logical circuits from end-customers terminate in Provider Edge (PE) routers. For each customer network, the router builds a Virtual Service Instance (VSI) to track locations within the customer subnet. This is analogous to the Virtual Routing and Forwarding (VRF) instances used in Layer 3 MPLS VPNs.

The PE router defines two MPLS labels to identify the traffic flowing from one customer and transport it over the network. An "inner" MPLS label identifies the individual PW associated with a given customer service instance. The PE router then adds a second label to the stack to identify the LSP through which the PW runs. The MPLS network switches the packet across the network based on the outer LSP label. At the remote PE router, the outer label is stripped off and the inner label is used to associate the incoming PW with the appropriate Attachment Circuit (AC) for transmission to the end-user location.

By using a two label stack, MPLS is able to provide secure VLAN separation for each carrier customer while also ensuring that the MPLS core has no knowledge of individual customer networks.

In addition, MPLS provides multiprotocol support since the core network forwards traffic based on carrier imposed MPLS labels, which allow encapsulation and emulation of many lower level network protocols (IP, Ethernet, PPP, ATM, Frame Relay, HDLC, etc.). Inherent in the MPLS label is also a set of EXP bits that are used to carry Quality of Service (QoS) data allowing routers to forward traffic based on user-specified traffic prioritization schemes. These concepts are further explored later in this white paper.

### Separation of Provider and Customer Control Plane

Unlike IP routing, MPLS divides the customer and carrier segments of the network at the control plane. Communication between the PE and P routers is based on a set of signaling protocols, such as LDP and RSVP, setting up LSPs over which customer traffic is forwarded based on MPLS labels stored in a Label Forwarding Information Base (LFIB). Extended L3 routing protocols such as OSPF-TE or ISIS-TE are used by the network to learn the MPLS network topology. This entire process is transparent to the customer.

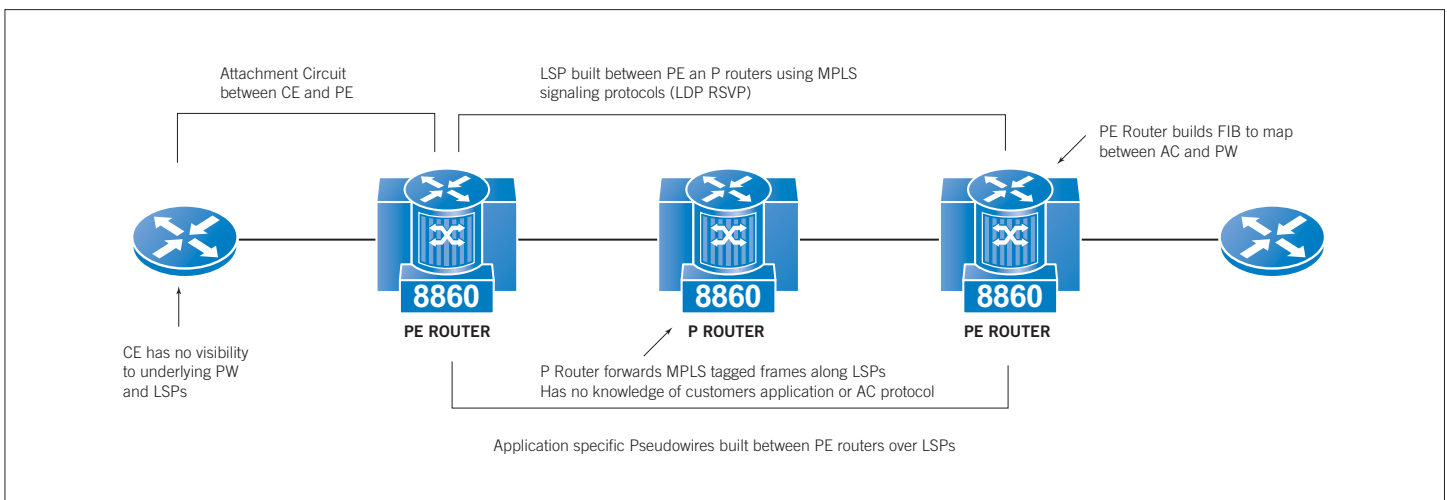


Figure 2. Control plane separation using MPLS LSPs and PWs

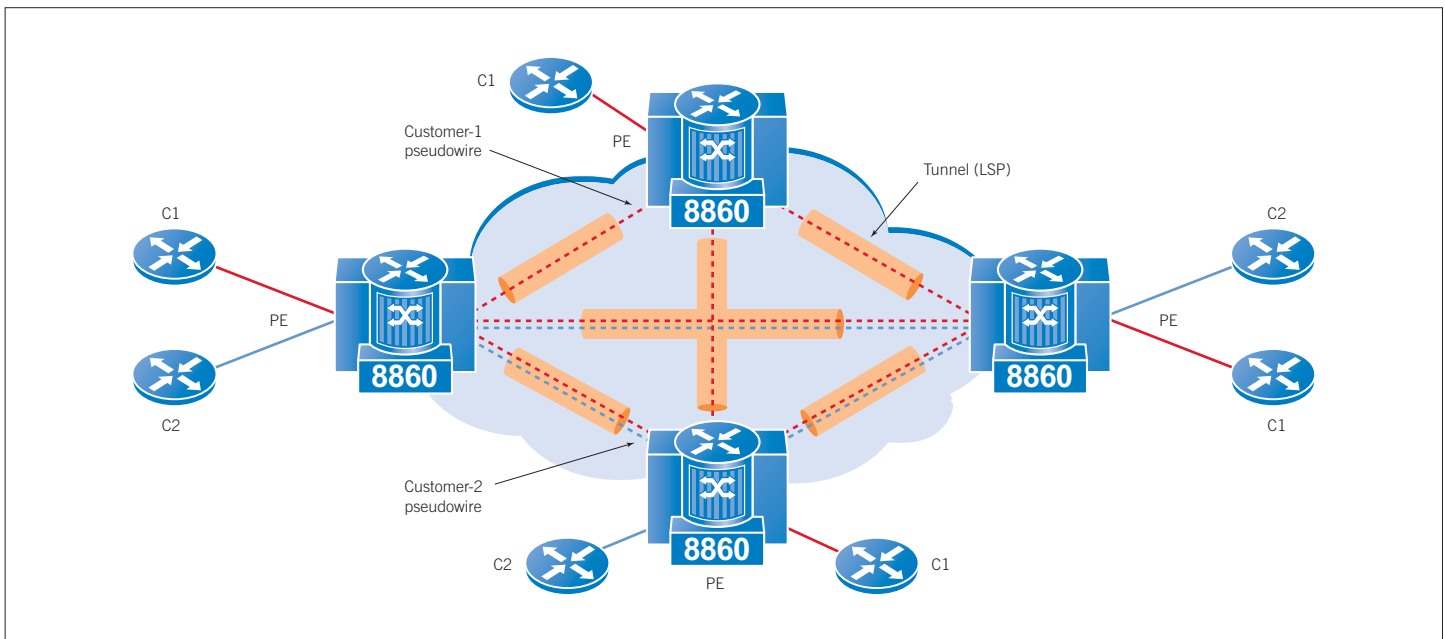


Figure 3: VPLS overview

Likewise the P routers are unaware of the details of the customer networks, which are restricted to the domain of the PE routers. PE routers receive customer traffic, build forwarding tables on a per-customer basis and label the packets based on the destination PE. This control plane separation, along with dual label stacking, allows the service provider to build a sophisticated multiprotocol, multiservice core network while presenting to each end-customer a much simpler interface into a private Ethernet VLAN. The control plane separation hides the inner complexity from the end-user.

### Pseudowires

In order to emulate an Ethernet switch over this MPLS core, a set of virtual circuits are built between PE routers to provide forwarding paths for the incoming Ethernet frames. These virtual circuits are termed pseudowires (PW).

Construction of a PW begins with the configuration of two uni-directional LSPs between PE routers on the network edge. Once this LSP tunnel is established, a PW is assigned to a pair of LSPs. As a result, all traffic engineering and resiliency features built into the LSP will also be in effect for the PW. Once service-specific PWs are built over the LSPs, packet forwarding across the network proceeds in a well-defined manner.

On the most basic level, PWs can be built and offered as standalone connections between a pair of customer sites. These Ethernet Private Line offers are also often referred to as Virtual Private Wire Service (VPWS) or Ethernet Virtual Private Line (EVPL) and are popular service provider offerings in their own right.

Although the focus of this white paper is Ethernet, PWs can also emulate FR, ATM, PPP, HDLC and a variety of Layer 1 TDM services for a true multiservice network. From the PE Router perspective, each customer connection, regardless of protocol, is defined as a CE-facing Attachment Circuit (AC).

### VPLS — An Ethernet Virtual Switch over the WAN

VPLS expands the concept of Ethernet service from a simple two-point circuit to a multi-port bridged service. VPLS is built over a full mesh of LSPs among all PE routers supporting the carrier service. A full mesh of Ethernet PWs is then provisioned over the LSP tunnels between all PE routers at which individual customer sites terminate. Customer CE routers or switches forward Ethernet frames as if they were attached to a local LAN Switch. P routers in the network core build LSPs and forward traffic with no awareness of the overlying application. PE routers stand at the carrier edge, maintaining MAC address tables for each customer network to associate individual traffic flows with the appropriate LSPs.

To establish this networking solution, VPLS performs several functions to emulate characteristics of an Ethernet switch:

#### Fully Meshed Loop Free Connectivity

Although MPLS networks are designed over an advanced topology of PE and P routers, a full mesh of PWs masks this underlying complexity to the VPLS instance. As a result, PE nodes will always see a direct link to remote sites, removing the need for the service provider to run Spanning Tree Protocol within the network core. The less complex Split Horizon ensures that the network remains loop free.

The lack of reliance on Spanning Tree Protocol also decreases the time required to re-establish connectivity in the case of a circuit outage since the spanning tree does not need to be rebuilt following a link failure. In the case that the customer does require use of Spanning Tree to provide options for “Back-door” failover connections around the VPLS network, Spanning Tree BPDUs can be tunneled through the network.

A variety of redundancy options are provided within the MPLS network to ensure that link or node outages within the network are quickly restored through service re-routing. These methods can provide sub-second link restoration times similar to those seen with SONET services. This subject is covered in more detail under the Advanced VPLS Features section of this white paper.

### MAC Address Learning

Ethernet transmits customer traffic based on MAC addresses. To emulate an Ethernet switch over an MPLS network, each carrier PE router maintains a Forwarding Information Base (FIB) of MAC addresses including MAC to PW mappings. A separate FIB is maintained for each VPLS service instance to provide traffic separation between customer networks.

VPLS maintains the MAC learning and bridging concepts but adapts them for the MAN/WAN. Reachability information is sent and received through traditional bridging in the data plane (ex. processes as defined in 802.1d and 802.1q). This contrasts to L3 MPLS VPNs in which the control plane discovers reachability information through traditional routing protocols. MAC address aging functions are also included in the VPLS standards to ensure that the FIB is up to date.

Replication of an Ethernet switch also requires support for flooding of unknown MAC addresses. The source router broadcasts the frame with an unknown address to all PE routers participating in the VPLS. The receiving router that owns the destination MAC responds. In a similar fashion, VPLS also supports Broadcast and Multicast traffic over the network.

### Connecting to VPLS — An Enterprise Customer Perspective

One of the primary benefits of VPLS service is the creation of a simple Ethernet network interface for end users. From a carrier standpoint, VPLS service requires extensive understanding of MPLS principles to design a traffic engineered core network of meshed LSPs providing a suite of network services. From the Enterprise customer standpoint, things look much different.

- Customer edge devices interface with VPLS service through simple and familiar Ethernet interfaces into a virtual LAN switch. From the customer standpoint, this appears similar to plugging into an on-site LAN switch.
- Since VPLS is a bridged Ethernet service, a customer can use either a LAN switch or a router to connect to the service. No Layer 3 routing protocols need be run to network together diverse locations.
- When routers are used to connect to the carrier network, VPLS requires no route sharing with the service provider. Customers maintain complete control over routing and their IP address space. There is no need to run BGP over the WAN, as is often seen in L3 MPLS VPNs. Standard Interior Gateway Protocols such as OSPF and RIP can be used over the entire network. As a Layer 2 network, the carrier service network has no interaction with the routing protocols.




Figure 4: Connecting to VPLS

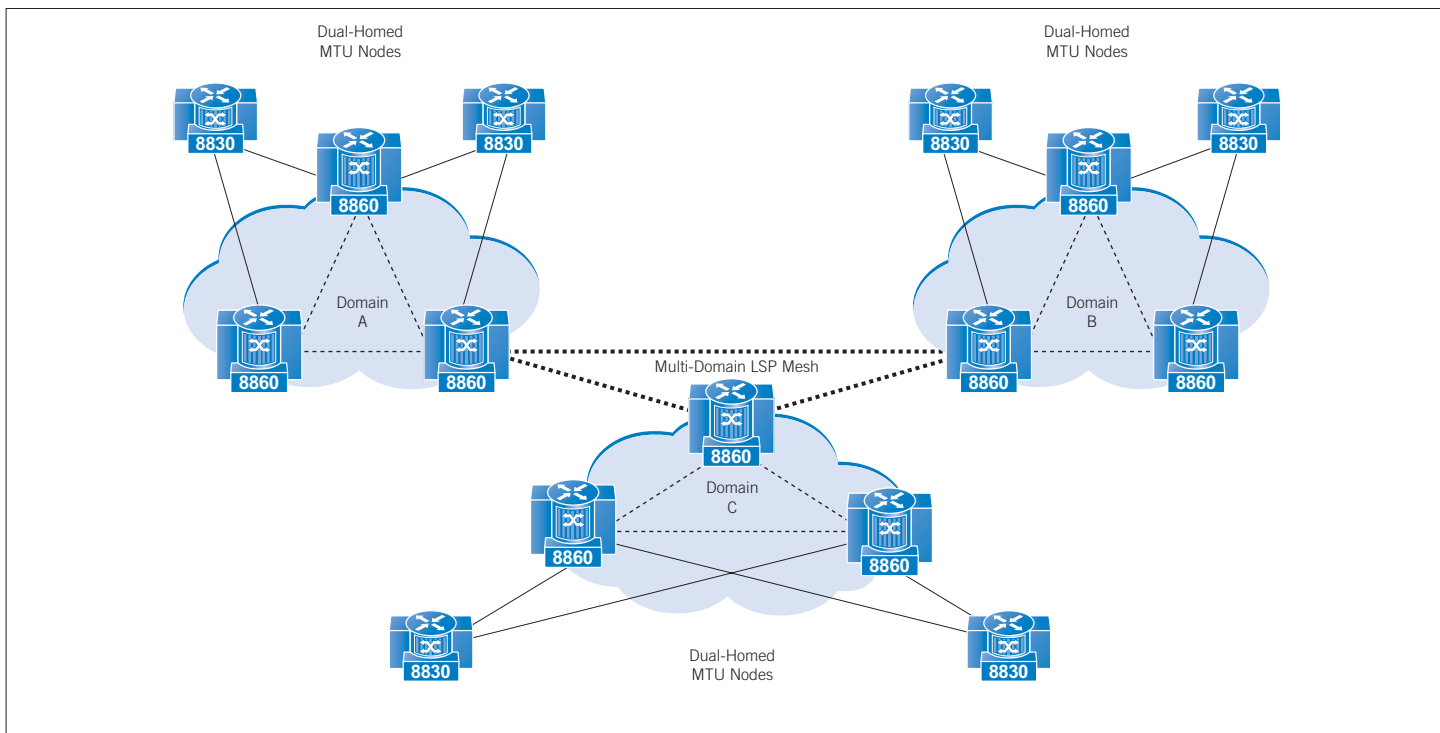


Figure 5: Hierarchical VPLS

- Since VPLS is a Layer 2 service, any type of Layer 3 protocol is supported over the networks. This benefits networks that continue to utilize legacy protocols such as IPX and SNA. Similarly, IPv6 packets can be forwarded over the network for those networks that do or will require migration to IPv6.

Unlike IP-SEC based VPNs, VPLS does not require complex customer site devices since MPLS provides a secure network infrastructure. Unlike L3 VPNs over MPLS, it does not impose routing requirements on the customer side in order to interact with the carrier network. Similar Ethernet-based offers have proven very popular in Metropolitan Area Networks. VPLS now offers the option to extend these successful network options into the national and global networks, adding extensive redundancy and QoS features to the Carrier Ethernet portfolio.

## Advanced VPLS Features

### Network Scalability

Carrier Ethernet services require a level of scalability that far exceeds traditional LAN-based Ethernet. A service provider may need to aggregate hundreds or even thousands of customer networks into one or two provider edge devices. Scalability features, which easily served the needs of single enterprises, will not serve the needs of a national or global carrier network. VPLS addresses these concerns with both H-VPLS and service instance scaling.

### Hierarchical VPLS (H-VPLS)

H-VPLS is a core feature of VPLS outlined in the original VPLS Standard (RFC 4762). It is designed specifically to support VPLS deployments in very large networks. H-VPLS combines two distinct network strategies: one to scale networks from the PE routers out to customer sites, and a second to provide scalability in the core network.

On the edge of the network, instead of building a full LSP mesh between all network locations, a remote carrier node designated as an MTU maintains simple point-to-point PW connections to the nearest PE routers. This results in a reduction of the number of routers involved in the LSP full mesh. This “two-tiered” attribute allows lower cost Ethernet devices running 802.1q VLAN or Q-in-Q to function as remote PE devices. H-VPLS further enhances network resiliency by making it possible to connect MTUs in a redundant fashion, with one active link and one standby link, connected to redundant PE routers using Spanning Tree Protocol to control routing loops.

Multidomain VPLS creates a third hierarchical tier in the network core. LSP growth will normally accelerate exponentially since VPLS requires a full-mesh of LSPs between PE routers. To control LSP growth, a network can be divided into several domains with a central inter-domain mesh. Border routers in each domain are used to deploy a full mesh of inter-domain spokes, as shown in Figure 4.

### Service Instance Scaling

In traditional bridged Transparent LAN services, 802.1q VLAN IDs are used to mark traffic into individual service instances. Using the enhanced 802.1ad standard, also referred to as Q-in-Q, two VLAN tags can be used to mark Ethernet frames — one VLAN tag used by the service provider to mark the service instance, and a second inner tag used by the customer for internal traffic separation. VLANs are numbered from 0 to 4095 using a 12 bit VLAN ID in the VLAN header. This allows the creation of 4096 separate customer instances across the service provider network, with each individual instance able to support up to 4096 customer defined VLANs for internal use.

While 4096 VLANs are enough to satisfy the needs of an individual enterprise, it can quickly become insufficient for carrier deployments. The Tellabs® 8800 Multiservice Router (MSR) Series addresses this challenge by allowing the service provider to identify each broadcast domain by a 15-character alphanumeric VLAN-ID. This feature improves carrier scaling to over a million possible VLANs per node in addition to making it much easier to assign and recognize customer VLANs during network installation and troubleshooting.

This scalability improvement extends into the MPLS backbone. When translating VLANs into the MPLS network, many routers translate the 802.1q VLAN ID directly to the MPLS header Virtual Circuit ID (VCID) field. Although the MPLS VCID is 20 bits long, it is effectively limited in scale to the 4096 VLANs possible with the VLAN header. Since the Tellabs 8800 MSR uses the alphanumeric VLAN ID, it can take full advantage of the scalability built into the MPLS standard.

### Resiliency

As networks move towards an MPLS-based infrastructure and service providers take advantage of new flexibilities in service offerings and backbone consolidation, they still must maintain the sub-second service restoration intervals, which customers expect from modern carrier networks. Redundancy and resiliency features in MPLS, and therefore VPLS networks, are available in a variety of forms:

- Link Aggregation allows multiple links to be bundled together to provide physical link redundancy on individual links.
- At the MPLS layer, primary and back-up LSPs can be set up to failover between a set of PE in the network.
- Fast Re-Route utilizes the concept of back-up LSPs to create pre-programmed restoration paths that can be used to route around either individual LSP link- or complete node-failures in SONET-like intervals.
- PW redundancy provides redundancy at the individual circuit level by failing-over a link from one remote PE to a second PE to allow routing around PE or Attachment Circuit failures.
- Bi-Directional Forwarding Detection (BFD) provides detection of failures directly to routing protocol clients for sub-second re-routing of individual customer circuits.

A combination of these methods, along with intelligent network architecture planning, creates a robust service offering. Additional OAM standards developed specifically for carrier Ethernet applications, such as 802.1ag service level OAM and 802.3ah link level OAM, provide additional tools to monitor and maintain VPLS networks.

### Interworking and Network Migration

Due to the underlying multiservice capabilities of MPLS, Ethernet WANs designed using VPLS service easily integrate with legacy FR and ATM networks. Carriers can extend the value of existing investments as well as offer a clean migration strategy to a new and more cost effective infrastructure.

Multiservice routers give the service provider additional methods to offer new integrated services over a single network. To ensure maximum operational flexibility and protection of existing investments, for example, Multiservice routers can extend VPLS service beyond the traditional reach of carrier Ethernet switches. Service providers can configure VPLS to support interworking with legacy FR and ATM networks, and to overlay VPLS on TDM and SONET infrastructures.

A customer connection into the VPLS network is defined as an Attachment Circuit, which need not be a native Ethernet link. For example, the customer connection could be an FR or ATM circuit, or Ethernet over an existing SONET physical circuit (ex. X.86 and GFP). This level of flexibility makes VPLS well-suited for integration with existing networks. It also simplifies the migration path from legacy networks to a new MPLS core without forcing CPE upgrades to support complex IP routing as can be the case when migrating to L3 MPLS VPNs. In addition, a multiservice router can support a variety of protocol interworking options so that entire ATM/FR networks can bridge directly into VPLS-based Ethernet services for new hybrid carrier product offerings. As illustrated in Figure 6, these advanced MSRs offer a clean path for enhancing and migrating legacy networks.

### Quality of Service

The original Ethernet standard was developed to run over an LAN where high speed connections to the desktop are the norm. It was not designed to support QoS for traffic prioritization. However, when running over a WAN where bandwidth is at a premium and latency emerges as a concern, QoS becomes a crucial network requirement, especially when running performance-sensitive applications such as voice and video over the data path.

IEEE 802.1p provides a method for marking Class of Service in Ethernet VLAN headers. VPLS takes advantage of these markings in running over an MPLS core. MPLS supports QoS with traffic prioritization marked in the EXP bits of the MPLS header. In a VPLS-based Ethernet VPN, an MSR translates the 802.1p bits from the VLAN header to the MPLS EXP bits to provide end-to-end QoS for Ethernet connections. In fact, VPLS traffic can also be prioritized based on a wide variety of header information, including both Layer 2 and Layer 3 Access Control Lists (ACLs). This blurs the lines between a traditional Ethernet implementation and higher level WAN technologies. Once again, while VPLS presents a simple Ethernet interface to the end customer, it actually provides a much more sophisticated service than traditional Ethernet.

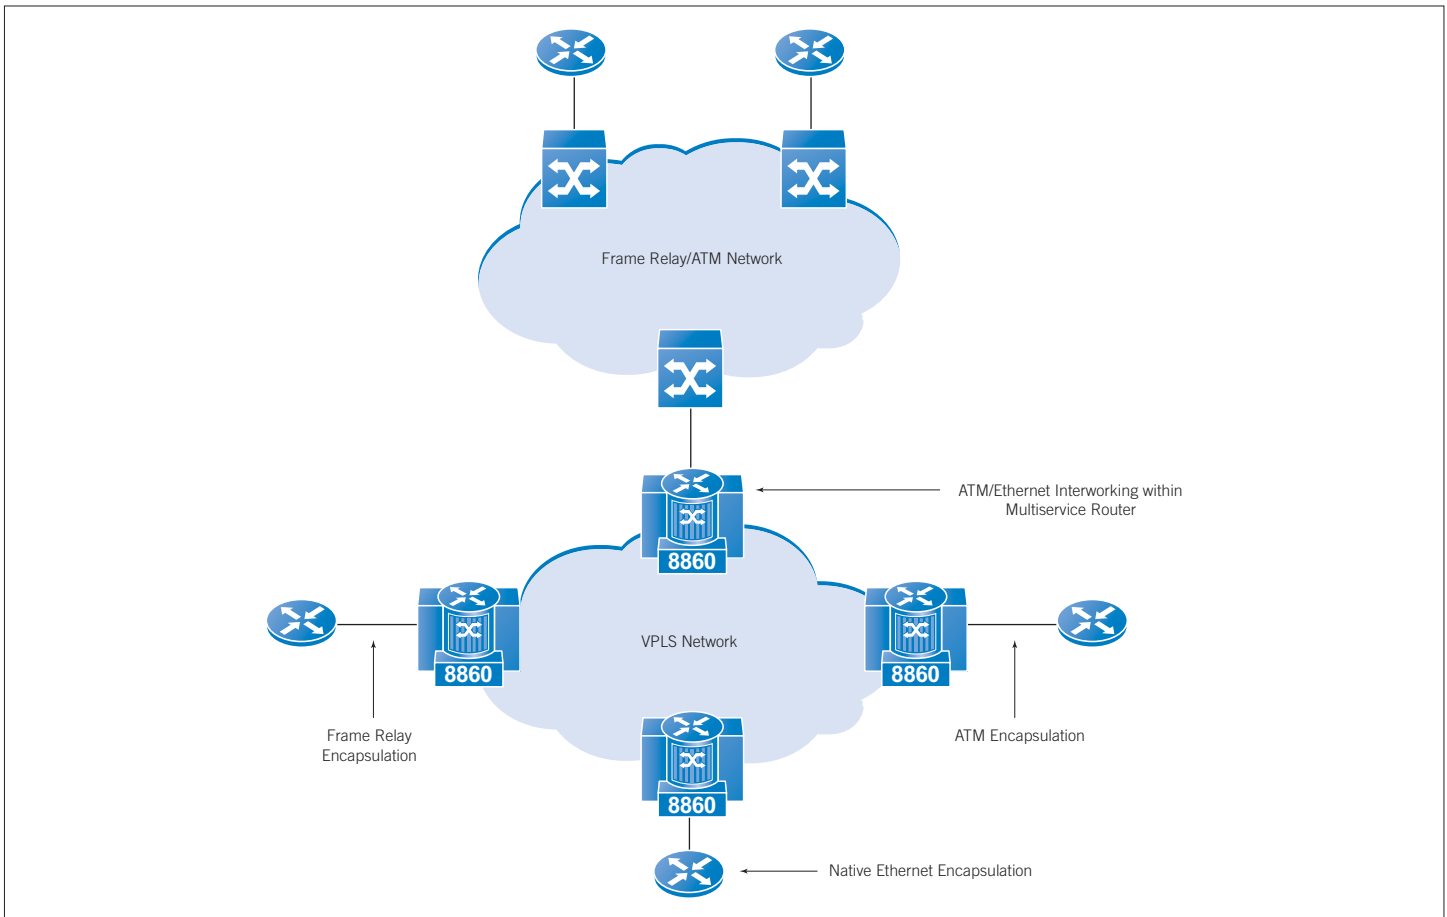


Figure 6: Legacy network interworking with VPLS using an MSR

## Conclusion — Ethernet Everywhere

VPLS creates benefits for both the service provider and the enterprise customer. In both cases, companies confront a networking environment centered on two primary communication protocols: Ethernet and IP. Ethernet has come to dominate the LAN environment. VPLS provides powerful tools to allow Ethernet to extend to the WAN as a sophisticated and scalable connectivity service.

On the end-user side of the network-demarkation line, businesses gain access to a new affordable method to link remote offices. In addition, because Ethernet operates at a Layer 2 level, enterprise organizations do not have to contend with the complexity of IP routing protocols or interoperating IP routing with carrier networks. Since VPLS provides an Ethernet handoff, routers are not required at Enterprise customer locations. End-users can continue to use the familiar Ethernet user-to-network interface or protect their existing

investments in ATM and FR access equipment and still converge their voice, video and data traffic onto a single, cost-effective, high-performance WAN.

For the service provider, VPLS creates significant new revenue streams, enabling providers to maximize the return on their MPLS investments in the core network. The scalability of VPLS allows service providers to cost-effectively deploy it to enterprise customers, all while controlling the Capital and Operating Expenditures associated with market growth. Finally, VPLS allows service providers to differentiate their Ethernet offerings in the competitive marketplace with guaranteed Service Level Agreements (SLA) to support critical enterprise applications that require more than best effort services.

VPLS provides a logical next step in the continuing evolution of Ethernet from a 10 Mbps shared LAN protocol to a multi-Gbps global service.

### North America

Tellabs  
One Tellabs Center  
1415 West Diehl Road  
Naperville, IL 60563  
U.S.A.  
+1 630 798 8800  
Fax: +1 630 798 2000

### Asia Pacific

Tellabs  
3 Anson Road  
#14-01 Springleaf Tower  
Singapore 079909  
Republic of Singapore  
+65 6215 6411  
Fax: +65 6215 6422

### Europe, Middle East & Africa

Tellabs  
Abbey Place  
24-28 Easton Street  
High Wycombe, Bucks  
HP11 INT  
United Kingdom  
+44 870 238 4700  
Fax: +44 870 238 4851

### Latin America & Caribbean

Tellabs  
1401 N.W. 136th Avenue  
Suite 202  
Sunrise, FL 33323  
U.S.A.  
+1 954 839 2800  
Fax: +1 954 839 2828

Statements herein may contain projections or other forward-looking statements regarding future events, products, features, technology and resulting commercial or technological benefits and advantages. These statements are for discussion purposes only, are subject to change and are not to be construed as instructions, product specifications, guarantees or warranties. Actual results may differ materially.

The following trademarks and service marks are owned by Tellabs Operations, Inc., or its affiliates in the United States and/or other countries: TELLABS®, TELLABS and T symbol®, and T symbol®.

Any other company or product names may be trademarks of their respective companies.

© 2008 Tellabs. All rights reserved.  
74.1932E Rev. A 3/08