MPLS in Private Networks

April 2005

By Absolute System Solutions

info@absolutesystem.jp

Introduction

The wide area network (WAN) brings indisputable value to organizations of all types. Through the WAN, an enterprise can instantaneously communicate between all of its locations and a wide variety of customers, suppliers and distributors. (Note: Throughout this document, the term "enterprise" refers to any organization, whether it is a corporation, government entity, or education system.)
Four key trends are driving the way that organizations design WAN infrastructures:

- **Pace of Business Change** – It is quite common for a company to go through mergers and acquisitions and to make frequent changes to its set of suppliers and distributors, as well as to gain new customers. As a result, the enterprise WAN needs to be extremely flexible and responsive to change.
- **Linkage to Business Goals** – The growing linkage between a robust network infrastructure and achieving business goals places new demands on the network for greater resiliency and scalability. To be considered resilient, a network has to quickly reroute traffic around a failed component - typically in 50 ms or less to preserve voice calls. The scalability of a network refers to factors such as the maximum speed of a WAN link or the maximum number of virtual circuits supported by the network.
- **Convergence of Network Infrastructure** – Enterprises are consolidating a wide variety of technologies (i.e., ATM, Frame Relay), protocols (i.e., IP, IPX, and SNA) and traffic types (i.e., data, voice, and video) onto a single network infrastructure. Supporting a single converged infrastructure is notably less costly than supporting multiple networks. However, a converged network infrastructure does introduce some significant challenges. In particular, organizations that deploy a converged infrastructure must ensure that the network can effectively and efficiently meet the demands of the disparate traffic types.
- **Traffic Isolation** – In the evolution of WAN design enterprises are looking to isolate traffic based on the organization responsible for the traffic. The isolation of traffic serves two purposes. It increases security by giving each organization within the extended enterprise access to only its own traffic. Isolation also increases network stability, since actions taken by a given entity will only affect that entity. The goal of this white paper is to detail an emerging approach to wide area networking that gives enterprises a path for evolving their network infrastructure to respond to these trends. The approach involves the deployment of Multi-Protocol Label Switching (MPLS). MPLS has been extensively deployed in service-provider backbones over the last few years. This paper will demonstrate the value of deploying MPLS within a private network.

MPLS Fundamentals

A router that supports MPLS-based forwarding is generally referred to as a label-switching router (LSR). It is common to refer to the first LSR in the data path as the ingress LSR, to the last LSR in the data path as the egress LSR, and to LSRs on the data path between these two as core LSRs. As the name implies, in an MPLS network each packet contains a label. A label is always 20 bits in length and is part of the 32-bit MPLS header. The label is assigned at the ingress LSR.
The forwarding function of a WAN is responsible for transporting a packet across the network, based on the information found in a routing table. The WAN control function is responsible for the construction and maintenance of the routing table, as well as for communicating routing

information to other nodes. One of the key attributes of MPLS is that it separates the forwarding and control functions. The separation of these two functions allows each function to be independent of the other. The MPLS control function uses a standard routing protocol such as OSPF to create and maintain a forwarding table. When a packet arrives at an LSR, the forwarding function uses information contained in the packet's header to search the forwarding table for a match. The LSR then assigns a label to the packet and forwards the packet to the next hop, in what is referred to as the label-switched path (LSP). All packets with the same label travel the same LSP from origin to destination. Unlike standard routing protocols, it is possible to have multiple active paths between two endpoints (Figure 1).

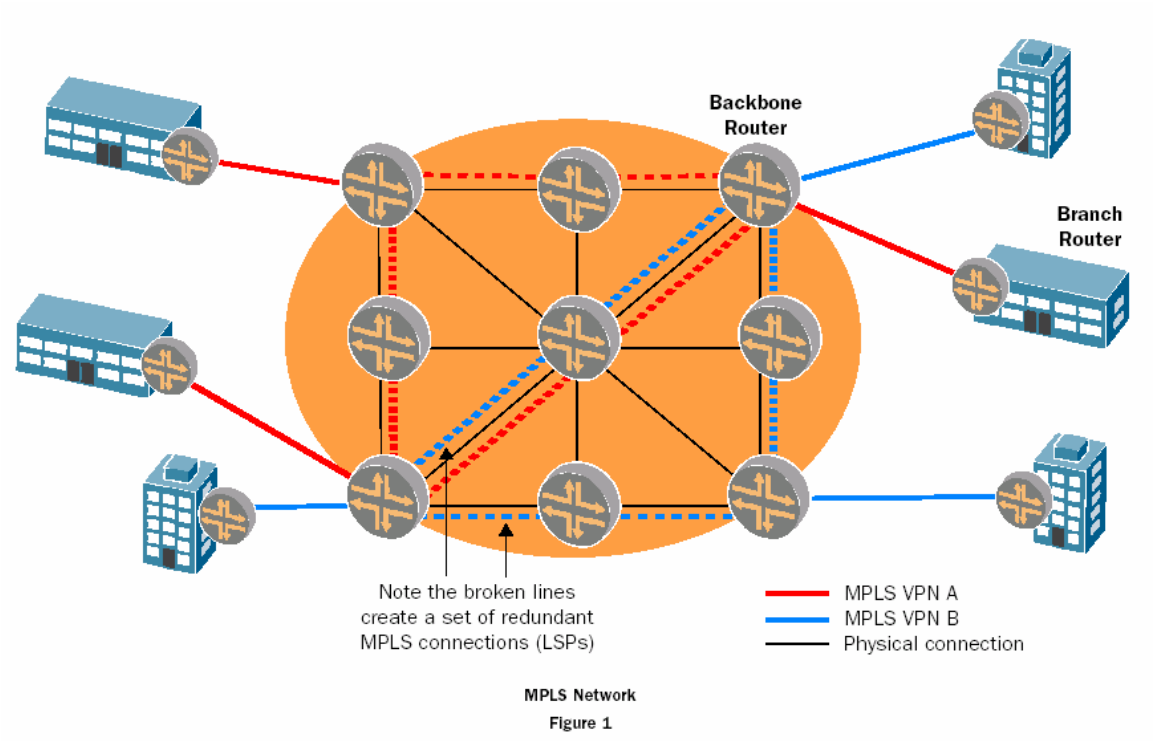

MPLS Network
Figure 1

The core LSRs ignore the packet's network layer header. Instead, when a packet arrives at one of these LSRs, the forwarding component in the LSR uses the input port number and the label to perform a search of the forwarding table. When a match is determined, the forwarding component replaces the label and directs the packet to the outbound interface for transmission to the next hop in the LSP.

The MPLS Value Proposition

For a company to achieve its business goals, an enterprise WAN must be highly scalable. An MPLS network is capable of supporting thousands of virtual private networks (VPNs) on a single physical infrastructure. Part of the reason that MPLS is capable of supporting so many VPNs is that its division of labor reduces the requirements for complete end-to-end peering across the network. Frame Relay and ATM are both Layer 2 (L2) technologies based on permanent virtual circuits (PVCs). Network organizations typically terminate these L2 PVCs in an IP router. This approach is highly complex, because it requires mapping between two architectures designed for fundamentally different functions. MPLS reduces network complexity in part because MPLS

integrates both L2 switching and Layer 3 (L3) routing in a single, uniform, standards-based protocol hierarchy. This section of the white paper will detail additional aspects of MPLS as an important enabling technology for implementing carrier-class networks. The phrase “carrier class network” refers to a private network that offers a range of services to its users, just as carriers offer multiple services to their customers. The phrase also refers to a network that is highly flexible and robust, and which enables the consolidation of technology and traffic types, as well as the isolation of traffic to meet many different needs.

Traffic Separation and Network Virtualization

Security is a top-of-mind issue for anyone concerned with deploying and managing a WAN. One of the reasons is the WAN’s central function of enabling communications with customers, suppliers, and distributors – all demanding assurances that the WAN connection cannot be used to launch a security attack against them. Another reason that WAN security is important is the recent shift in how enterprises think about security. In the past the focus was on providing security only at the perimeter of the network. Organizations took this approach because it was assumed that the vast majority of security attacks came from outside the enterprise.

Today it is widely accepted that the majority of security incidents originate from within the enterprise. So in addition to providing separation between the communications of an enterprise and its external customers and partners, the WAN must now also keep separation between the communications of individual departments and work groups.

Furthermore, in many large enterprises there is a growing requirement to give each organization full domain and control over their network. This approach facilitates autonomous business operations, as well as the fluid entry and exit of different groups.

In a virtual private network (VPN), multiple traffic streams run over a common infrastructure in a way that each traffic stream appears to be running over a private network. The ability to implement VPNs is a key requirement of converged networks. By implementing VPNs, a network organization can assign unique security and quality of service (QoS) parameters to each traffic stream, while enabling more autonomous business operations.

MPLS enables the deployment of VPNs by supporting a simple, flexible, and powerful tunneling mechanism. An enterprise network organization can deploy a VPN by provisioning a set of LSRs to provide connectivity among the sites that comprise the VPN. Each ingress LSR places traffic into an LSP based on the combination of the packet’s destination address and its VPN membership information.

One of the security mechanisms that is inherent in MPLS-based VPNs is traffic separation. In order to separate traffic, each MPLS-enabled VPN is assigned to a unique virtual routing and forwarding (VRF) instance. Traffic destined for each VRF carries its own label value, so each VPN is kept logically separate from every other VPN. Utilizing these techniques, MPLS based VPNs offer the same level of logical security as ATM or Frame Relay virtual circuits, with the added advantage of being delivered on a single converged network Each organizational entity can operate its own VPN, setting up its own networking policies and IP addressing schemes.

MPLS also provides encapsulation mechanisms to carry any legacy or proprietary, non-IP protocols being used by the disparate, independent groups.

Traffic Consolidation and MPLS Service Classes

WAN design has always taken into account the need to provide acceptable levels of availability, delay, jitter, and packet loss. However, the definition of what level is acceptable changes markedly when companies consolidate many traffic types onto the WAN. For example, many users are accustomed to not being able to access their email when their server is down. However, these same users expect that the voice network will be available 100% of the time.
A converged network must provide far more stringent levels of performance for at least some of the traffic than is necessary for most data applications. The ITU (International Telecommunication Union) recommends that the one-way, end-to-end delay associated with a voice call does not exceed 150 ms. In addition, jitter should not exceed 40 ms and packet loss should not exceed 0.5%. Implementing MPLS gives an enterprise network organization tremendous flexibility in how it assigns packets to LSPs. The assignment can be based on a combination of factors, such as the source address, the destination address, the application type, the point of entry into (or exit from) the MPLS network, as well as class-of-service (CoS) information.
As a result, a network organization can take any type of user traffic and associate it with an LSP designed to satisfy its specific requirements. For example, the network organization could establish four classes of traffic:
• Voice
• Video
• High-priority data
• Low-priority data

Each of these classes of traffic is mapped to an LSP that has been designed to meet the required QoS. In the case of voice traffic, the LSP would be designed to provide a level of delay, jitter, and packet loss that is in line with the parameters previously discussed.
There are two approaches that an enterprise network organization can take to implement MPLS service classes. In one approach, there is a single LSP between a pair of edge LSRs. Traffic that flows on that LSP is placed into a queue on the LSR's outbound interface, based on the precedence bits in the MPLS header as set by the application. DiffServ-Aware Traffic Engineering (DS-TE) is a technique that allows MPLS to enforce the level of priority requested by the application. In the second approach, there are multiple LSPs between each pair of LSRs. Each LSP can be traffic engineered to provide appropriate network parameters. For example, the ingress LSR could separate voice, video, high-priority data, and low priority data into their own LSPs.

Traffic Engineering and Fast Reroute

The typical enterprise WAN is comprised of IP routers interconnected by Frame Relay or ATM PVCs. In this type of WAN, the network organization has little control over how the traffic is routed. Traffic routing is controlled by a routing protocol such as OSPF, and it is likely that the packets will encounter congestion as they traverse the network. The result of encountering congestion is that the packet flow will experience significant jitter, and possibly packet loss. While moderate amounts of jitter and packet loss are acceptable for most data applications, they are not acceptable for voice and some collaboration applications.

Traffic engineering refers to the process of selecting the paths that traffic will transit through the network. Traffic engineering can be used to accomplish a number of goals. For example, a network organization could traffic engineer their network to ensure that no links or routers are over-utilized or underutilized. Alternatively, a network organization could use traffic engineering to control the path taken by voice packets to ensure appropriate levels of delay, jitter, and packet loss. Unlike traditional routing approaches, MPLS supports traffic engineering. MPLS-based traffic engineering allows network organizations to associate an LSP with whatever physical path they choose. MPLS also supports constraint-based routing, which ensures that an LSP can meet specific performance requirements before it is configured. In addition, tools that work on a per-LSP basis allow network organizations to identify utilization levels and plan accordingly.
MPLS-based traffic engineering also supports the rerouting of traffic around a failed link or router quickly enough that users of the network are not adversely affected. To achieve this fast restoration time, a backup LSP can be established at each node. The failover mechanisms are triggered by physical link or routing events that indicate that the link or node is down.
The traffic can be switched immediately to this backup LSP once the failure has been detected. MPLS with this capability can reroute traffic in under 50 ms, which is similar to SONET/SDH networks that carry the public telephony network.

VPN Deployment Options

An enterprise network organization has several options for deploying a private MPLS-based VPN. For example, MPLS-based VPNs can be deployed at L3 as well as at L2, and there are multiple choices for the implementation of a L2 MPLS-based VPN.
An alternative to building a purely private network is to deploy a private MPLS-based network and connect it to public MPLS services, or even to extend MPLS to smaller sites through GRE/IPSec tunnels over IP services. The range of available deployment options allows the IT organization to deploy the MPLS-based VPN solution that best meets its needs.

Layer 3 VPNs

The implementation of L3 MPLS-based VPNs is typically based on IETF RFC 2547bis. This class of VPN transports traffic across the network through the use of MPLS tunnels and Multiprotocol Border Gateway Protocol (MP-BGP) signaling (Figure2). Note that in Figure 2, BB refers to a backbone router that is running MPLS, while BO refers to a branch office router that is not running MPLS. This is the most common way that MPLS-based VPNs are currently deployed. However, an enterprise could also run MPLS in their branch office routers to extend its benefits to the network endpoints. Also note that each BB has a routing instance per VPN, known as a VRF.

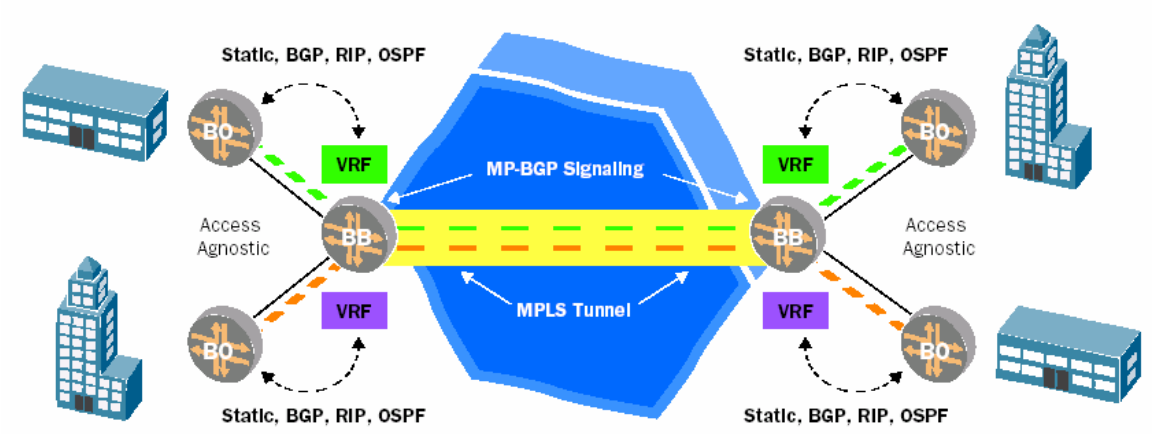

Layer 3 VPNs
Figure 2

Two of the advantages of MPLS-based L3 VPNs are that they are standards-based and easy to provision. This type of VPN also supports a wide range of access types and a variety of topologies, including full mesh, partial mesh, and hub and spoke.

Layer 2 VPNs

L2 VPNs, such as ones based on Frame Relay and ATM, are extremely common and are inherently multiprotocol. Realizing the importance of L2 VPNs, the IETF L2VPN working group has defined both encapsulation and label-distribution mechanisms that enable transporting non-IP protocols across an MPLS core network. Due to the multiprotocol nature of L2 VPNs, an L2 MPLS-based VPN presents an easy transition step for organizations that currently run legacy protocols but intend to migrate to an all-IP network over time. For example, a network organization can set up an L2 VPN to transport legacy protocols such as IPX or SNA over the core MPLS network, without having to encapsulate them inside of IP. One of the key features of an L2 MPLS-based VPN is the ability to create a tunnel as a LSP (Figure 3). Another key feature is the ability to use control protocols such as MPLS's Label Distribution Protocol (LDP) or BGP to set up emulated Virtual Circuits.

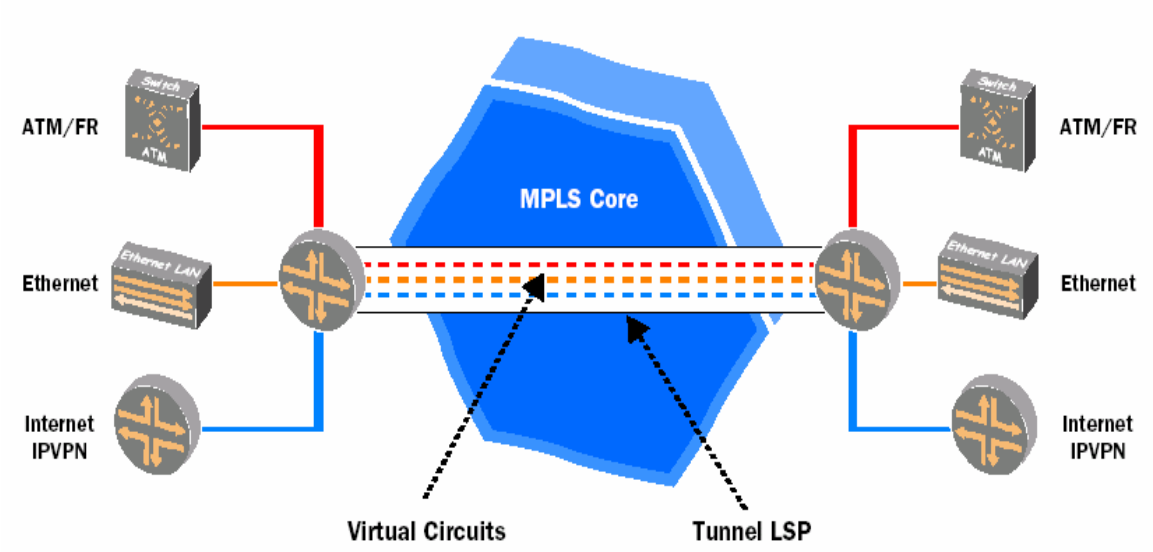

Layer 2 MPLS-Based VPN

Figure 3

Layer 2 MPLS-based VPN: Draft-Martini

The class of VPN called Draft-Martini refers to a set of Internet drafts that define L2 encapsulation and transport signaling methods. The approach is also called Pseudo Wire Emulation, because it is predicated on the construction of point-to-point circuits, that is, pseudo wires, over an MPLS core. A benefit of a Draft-Martini VPN over L3 options is that it can support a wide range of encapsulations, including Ethernet,
Frame Relay, ATM, High-Level Data Link Control (HDLC), and Point-to-Point Protocol (PPP). However, a Draft-Martini VPN does not scale well because each pseudo wire must be configured individually. A Draft-Martini VPN also introduces an additional protocol into the network, the Label Distribution Protocol (LDP), which is used to set up tunnels and to distribute labels.

Layer 2 MPLS-based VPN: Draft-Kompella

Draft-Kompella VPNs were designed to overcome the limitations of Draft-Martini VPNs. In particular, a Draft-Kompella VPN uses BGP instead of LDP for signaling. Since BGP is already used for routing across the MPLS core, a Draft-Kompella VPN does not require the introduction of an additional protocol in the network. There are additional benefits to Draft-Kompella VPNs. Through the use of auto-provisioning, they require minimal manual configuration. Also, through the use of Draft-Martini and other encapsulation techniques, a Draft-Kompella VPN supports a wide range of encapsulations.

VPLS Networks

A virtual private LAN service (VPLS) delivers a multipoint-to-multipoint Ethernet connection that spans one or more metropolitan areas (Figure 4). A VPLS provides connectivity between multiple sites over an existing MPLS backbone as if those sites were attached to the same Ethernet LAN.

One of the advantages of a VPLS is that sites connect to it using an Ethernet interface and an Ethernet domain is created between sites. Traffic entering on any of the Ethernet ports in that domain is then replicated to other ports in that domain as appropriate. To the end user the VPLS looks like a large distributed Ethernet segment.

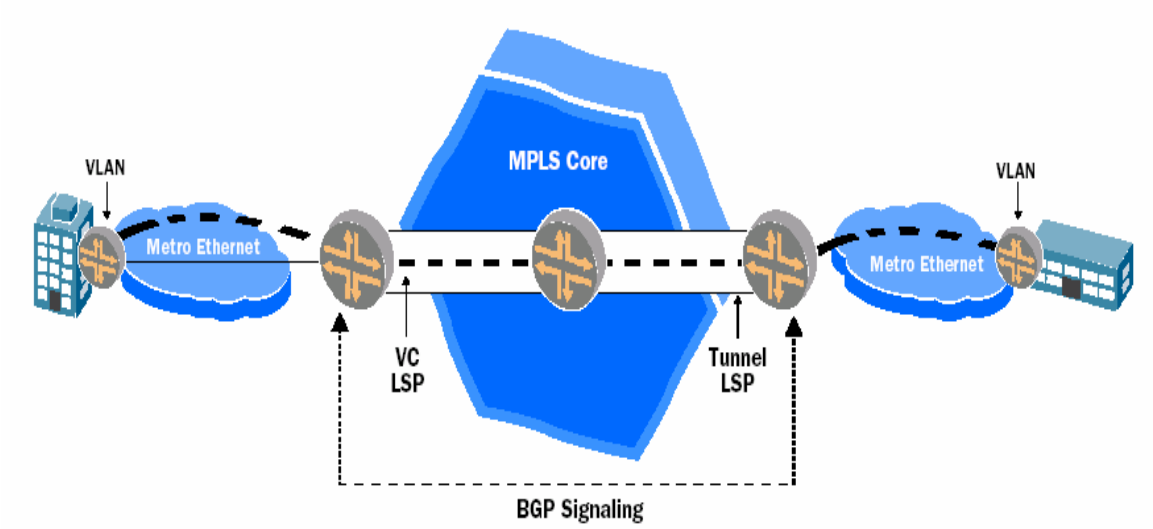

Virtual Private LAN Service (VPLS)

Figure 4

There are two proposed standards for implementing a VPLS. They differ based on their approach to the following:

• Auto-Discovery – What technique is used to enable backbone routers that participate in a VPLS domain to find each other?

• Signaling – What protocol is used to set up MPLS tunnels and distribute labels?

Draft Lasserre-Vkompella VPLS

This solution uses LDP for signaling and does not use a protocol for auto-discovery. Any network organization that implements it would have to know what backbone routers were a part of a VPLS instance. For every VPLS instance on a backbone router, the network organization would have to configure that backbone router with the addresses of all of the other backbone routers that are part of that VPLS instance. This approach is both operationally demanding and error prone, and it introduces another protocol (LDP) into the network.

Draft Kompella VPLS

A Draft Kompella VPLS uses BGP for both auto-discovery and signaling. Using BGP for auto-discovery greatly simplifies the configuration of VPLS without introducing an additional protocol into the network.

Hybrid MPLS Networks

Enterprise network organizations do not have to choose between implementing a private MPLS network and acquiring MPLS services from a provider. Rather, these organizations can

implement a hybrid network: a private MPLS network that is connected to one or more public MPLS networks. There are multiple ways to connect private and public MPLS networks. One approach is to connect at the 'VRF level' to maintain VPN connectivity across autonomous boundaries, while delivering full connectivity and reachability. This is a good approach for connecting peer networks. A second approach uses a central MPLS network to connect distributed islands of MPLS networks. This approach, often called the 'carrier's carrier' model, allows organizations with geographically dispersed locations to connect islands of remote MPLS networks together in a secure manner.
Additionally, where public MPLS networks are not available, companies can extend MPLS to branch sites through GRE/IPSec tunnels over any public IP service. A tunneled MPLS connection is also a cost-effective option for branch offices requiring backup links. Further, if the backup is over a broadband service, enterprises may choose to traffic engineer certain high volume data applications to always use the broadband connection.

Example Scenarios of Private MPLS Networks

This section of the document describes two scenarios that represent how enterprises could deploy a private MPLS network. In the first scenario, a financial organization operates a single network to connect a number of subsidiaries, all of which require full intra-subsidiary connectivity, but only very occasional inter0subsidiary connectivity. The subsidiaries have widely differing network needs. Some subsidiaries require only best-effort email service, while others need highly available access to time-sensitive transactional applications as well as VoIP support between locations. The solution for this organization is an MPLS network that utilizes L3 VPN technology, combined with traffic engineering and COS per VPN to meet the respective needs of the subsidiaries (Figure 5). Inter-subsidiary communication can be accomplished via route-leakage between the appropriate VPNs or via a firewall-based solution for tighter per-user level access control.

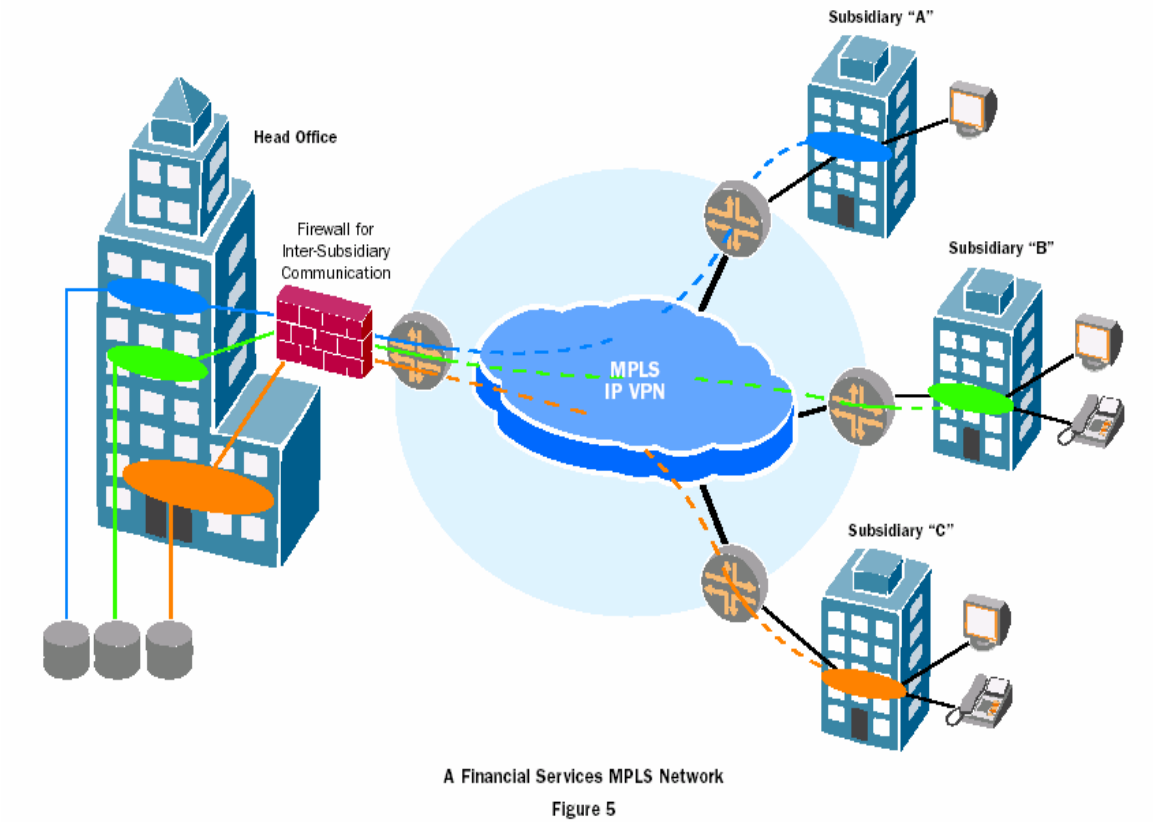

A Financial Services MPLS Network

Figure 5

In the second scenario an enterprise owns and operates a single network to support departmental and remote office connectivity to a number of key applications. The organization wants a scalable way of supporting the following:

• **Logical Separation of Departmental Traffic** – VLANs separate departmental traffic throughout the LAN infrastructure, and they want to maintain this logical separation across their WAN. This design requirement indicates the importance of security to the firm.

• **VoIP Deployment** – The organization is in the process of deploying VoIP to all of its facilities and places a premium on uptime - to the point that it has deployed dual connections into a number of remote and branch offices.

• **Access to Time-Sensitive Transactional Application** – Consistency of response time and performance are paramount, in part because a number of the organization's remote and branch offices are part of a distributed call center based on centrex services.

The solution is a private MPLS network with L3 VPNs deployed out to the branch and remote offices (Figure 6). The voice and data VLANs are mapped into VRFs on the branch office router and then transported across the WAN to other remote sites. The various VPNs can be traffic engineered to meet the needs of the organization's voice and data traffic. Given the organization's concern for security, IPSec could be implemented within this solution. In addition, the local router can be configured so that if one of the main links fails, all traffic can be rerouted within 50ms to the alternate path to ensure continuity of all user sessions.

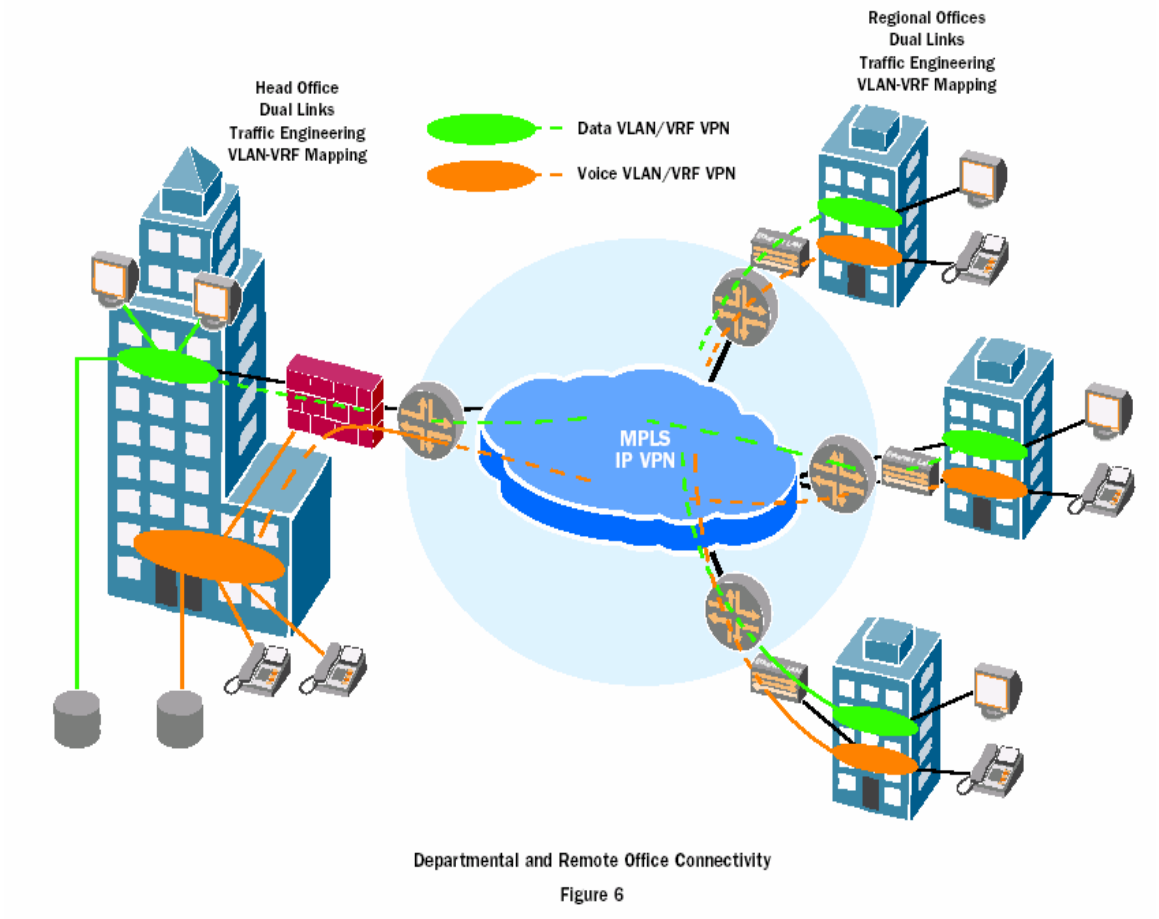

Departmental and Remote Office Connectivity

Figure 6

Call to Action

Private MPLS-based networks are not appropriate for every situation. They are, however, appropriate for enterprise network organizations with the following goals:

- To have more control over their network infrastructure
- To provide better performance, reliability, and efficiency
- To offer multiple classes of services to their user base
- To securely extend a virtual piece of the backbone network to organizational entities
- To ensure the performance of demanding applications
- To support the convergence of multiple technologies and/or multiple traffic types onto a single network
- To gracefully support legacy protocols while migrating to an all IP network.

Network organizations with these goals need to choose their hardware vendor carefully, based on a variety of criteria. For example, given that the MPLS-based network is intended to support all of the company's communications, network organizations should only choose a vendor whose products exhibit a track record of high performance, reliability, and scalability. In addition, it is critical that the chosen hardware vendor embrace the use of MPLS in enterprise networks. There are two key measures of whether a vendor embraces MPLS. One of these measures is the breadth

of the vendor's development of MPLS throughout their product line. The second measure is the depth of the vendor's development. For example, if an enterprise network organization intends to support video distribution over the network, it is critical that the vendor's MPLS implementation support point-to-multipoint connections. Companies that deploy an MPLS network will find that they are in a strong position to address the four trends discussed in the introduction to this paper. In addition, these companies will discover that the wide range of MPLS deployment options ensures that a private MPLS-based network can support virtually any set of complex business requirements.