IT INNOVATION REPORT | MARCH 2006

Innovation in MPLS-Based Services

By Jim Metzler

Kubernan
Guiding Innovation

Innovation in MPLS-Based Services

Introduction

MPLS (Multi-Protocol Label Switching) has garnered a lot of attention over the last few years. In particular, over that time frame the vast majority of the major carriers have implemented MPLS within their backbone networks. While there has been some discussion on the topic in the trade magazines, few carriers currently offer a service that extends MPLS all the way to the customer premise. There is also a nascent trend to have large IT organizations implement MPLS within their own private backbone networks.

This report provides the reader with a description of Layer 3 and Layer 2 MPLS-based Virtual Private Networks (VPNs). In addition, this document provides an assessment of the Layer 3 MPLS based VPNs that are available from three major carriers. Those carriers are AT&T, Sprint, as well as Verizon Business, which was recently formed by the merger of MCI and Verizon Enterprise Solutions Group. This document focuses on the innovation that is contained within those Layer 3 service offerings. None of these service providers currently offers a Layer 2 MPLS-based service, although that situation may change in 2006.

Throughout this document, Verizon Business, AT&T and Sprint will be referred to as The Carriers. The assessment of VPN services that is contained in this report was based in part on a set of questions that was sent to each of The Carriers.

The Motivation for Layer 3 MPLS-Based VPNs

Most large carriers are investing heavily in MPLS as a unifying network core technology that can support both legacy Layer 2 access services (Frame Relay and ATM) and emerging Layer 3 packet-based services. For most of these carriers, Layer 3 MPLS VPNs based on IETF RFC 2547bis represent both a revenue-generating service as well as a foundation upon which additional revenue-generating services may be based.

The reasons that are motivating enterprises to implement a Layer 3 MPLS-based VPN are the ability to:

IT Innovation Report

Published By
Kubernan
www.Kubernan.com

Cofounders
Jim Metzler
jim@ashtonmetzler.com

Steven Taylor
taylor@webtorials.com

Design/Layout Artist
Debi Vozikis
dvozikis@rcn.com

Copyright © 2006
Kubernan

For Editorial and Sponsorship Information
Contact Jim Metzler
or Steven Taylor

Kubernan is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.

Professional Opinions Disclaimer
All information presented and opinions expressed in this IT Innovation Report represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

- Combine multiple disparate networks onto a single network infrastructure
- Support any-to-any applications such as VoIP (Voice over IP)
- Enable efficient disaster recovery
- Prioritize applications in an easier fashion than is possible with Frame Relay or ATM
- Migrate off of legacy technologies such as Frame Relay and ATM in a seamless fashion
- Perform moves, adds and changes in an easier fashion
- Provide efficient access to multiple data centers

Fundamentals of Layer 3 MPLS-Based VPNs

A router that supports MPLS-based forwarding is generally referred to as a Label-Switching Router (LSR). It is also 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 communicating routing information to other nodes.

The MPLS control function uses a standard routing protocol such as OSPF to create and maintain a forwarding table.

MPLS Network
Figure 1

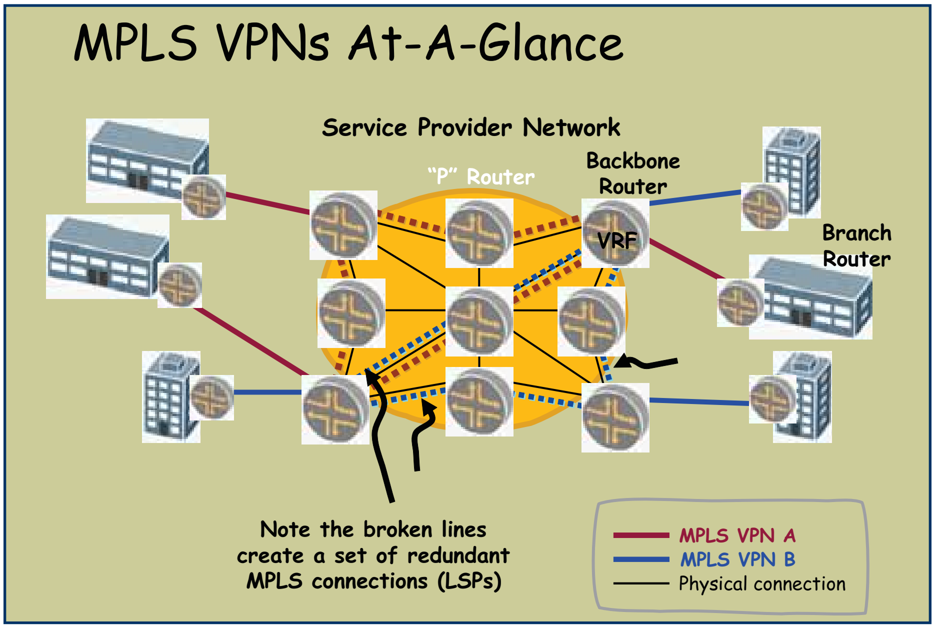

When a packet arrives at an LSR, the forwarding function uses information contained in the packet's header to search the forwarding table for a match. The LSR then assigns a label to the packet and forwards the packet to the next hop in what is referred to as the Label-Switched Path (LSP). All packets with the same label travel the same LSP from origin to destination. Also as shown in Figure 1, unlike the situation with standard routing protocols, it is possible to have multiple active paths from origin to destination.

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.

As previously mentioned, the dominant use of MPLS today is within service provider backbones. As a result, a lot of the MPLS-oriented terminology reflects this type of deployment. For example, the phrase *customer edge (CE)* *routers* is used to refer to IP routers at the customer site that provide Layer 3 connectivity to the PE (provider edge) router at the Service Provider's PoP (Point of Presence). The PE routers are LSRs that perform the MPLS encapsulation at the edge of the Service Provider's network, while the Provider (P) core routers are LSRs that simply switch the traffic to its destination.

The CE router may be configured with a static route to the PE or may exchange routing information with the PE via a routing protocol; i.e., RIPV2, OSPF, EIGRP, or eBGP. The CE is typically the management boundary between the service provider network and the customer network.

The support of routing protocols on the CE to PE boundary is one way the service providers differentiate their service offerings. With that in mind, The Carriers were asked to indicate which techniques they used to establish routes between the CE and the PE. Their responses are contained in Table 1.

On the PE router, a VRF (VPN Routing and Forwarding) table is configured for each separate VPN. This VRF virtual routing instance is the basic building block of an MPLS Layer 3 VPN. Routes learned from the attached CE router are populated into the VRF. The entry in the VRF table includes the original IPv4 route with a pre-pended RD (route descriptor). The ingress PE distributes the routes in the VRF using Multi-protocol BGP (MP-BGP). An attribute called the Route Target in the VRF table determines which PE routers in the Service Provider's network participate in the VPN and therefore need to receive BGP route distributions. When forwarding traffic, the ingress PE places two labels on the MPLS label stack. One label specifies the egress PE and the other label specifies the next hop P router

The core P routers are aware of only the PE routes, not of the VPN itself. Each P router on the path strips off the outer label and replaces it with a label for the next hop P router. When the final P router is reached, the outer label doesn't need replacement and the packet is forwarded to the egress PE that removes the label and forwards a native IPv4 packet to the egress CE.

There has been some discussion recently in the trade magazines about extending MPLS out to the customer premise. The Carriers were asked whether or not they supported such a service. Below are their responses:

1. Verizon Business

   They do not offer such a service, but would be open to it if the demand develops.

2. AT&T

   AT&T stated that they offer service to extend MPLS to the premise with and without managed CPE.

3. Sprint

   They do not offer such a service today but it is on the roadmap of things to support.

The MPLS Value Proposition

Traffic Engineering

The typical enterprise WAN is comprised of IP routers that interconnect either Frame Relay or ATM PVCs. In this type of a network, no organization has direct control over how the traffic is routed. The routing of the traffic is controlled by a routing protocol such as OSPF. That leads to a situation in which it is likely that as the packets traverse the network, they will encounter congestion. The result of encountering congestion is that the packets will experience significant jitter, and possibly packet loss.

Part of the innovation that MPLS offers is the ability to do traffic engineering. Traffic engineering refers to the process of selecting the paths that data traffic will transit through the network. MPLS-based traffic engineering allows an organization to associate an LSP with whatever physical path they choose. MPLS also supports constraint-based routing that ensures that an LSP can meet specific performance requirements before it is configured.

MPLS-based traffic engineering also supports the rerouting of traffic around a failed link or router quickly enough so as to not adversely affect the users of the network. To achieve this fast restoration time, a backup LSP can be established at each node. The fail-over 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 LSP once the failure has been detected.

Traffic Consolidation and MPLS Service Classes

The use of MPLS gives a service provider tremendous flexibility in terms of how it assigns packets to LSPs. This 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, the service provider can take any type of user traffic and associate that traffic with an LSP that has been designed to satisfy the requirements of that traffic. For example, the carrier may establish four classes of traffic. For the sake of example, those classes could be:

- Real Time
- Video
- Business Critical
- Best Effort

There are two approaches that a service provider can take relative to implementing 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 each LSR's outbound interface based on the setting of the precedence bits in the MPLS header.

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 put real-time, video, business-critical and best effort traffic each in its own LSP.

The Carriers could have differentiated their service offerings by choosing different approaches to implementing MPLS service classes. However, each of The Carriers implement traffic engineering the same way – by carrying multiple classes of traffic on the same trunk.

Service Classes Implemented by The Carriers

This section will describe the service classes implemented by The Carriers. As will be seen, there are significant differences in the approach that The Carriers take to implementing service classes. With virtually no exceptions, the following sections contain the exact responses that The Carriers gave when asked to describe their service classes.

Verizon Business

Verizon Business's Enhanced Traffic Management (ETM) is a CoS offering that allows customers to assign five traffic priority classes with up to eight priority levels. Real-Time traffic priority adds jitter as an SLA parameter. The five classes supported are:

- Expedited Forwarding (EF) is dedicated for real-time applications such as voice. The Gold CAR (Committed Access Rate) is assigned to the EF class. Traffic marked EF has the highest traffic priority. Any traffic that exceeds the subscribed EF/Gold CAR is dropped.
- Assured Forwarding 4 (AF4) is used for either video or business critical applications such as SAP, Siebel, PeopleSoft, or Point of Sale (POS).
- Assured Forwarding 3 (AF3) is associated with business critical applications, i.e., SAP, Siebel, PeopleSoft, POS, TN3270 emulations, Citrix. The primary difference between AF4 and AF3 is the AF4 class is associated primarily with video applications.
- Assured Forwarding 2 (AF2) is ideal for Telnet, Extranet Web Applications, general data applications.

- Best Effort (BE) has the lowest forwarding priority and is typically used for FTP, Database Synchronization, e-mail, web surfing. Traffic marked BE has the lowest priority.

AT&T

As a default, AT&T implements four data classes with two classes supporting further sub-categorization of service. AT&T stated that other COS models can be implemented on a custom basis.

AT&T's default service classes are:

- **CoS 1 –** This class is indicated with DSCP Expedited Forwarding (EF) and is intended for real-time applications such as interactive voice or video.
- **CoS2 (In Contract) –** This class is indicated with DSCP Assured Forwarding 31 (AF31) and is intended for time sensitive, mission critical, low bandwidth, bursty data applications.
- **CoS2 (Out of Contract) –** This class is indicated with DSCP Assured Forwarding 32 (AF32) and is intended for time sensitive, low bandwidth, bursty data applications. CoS2/InContract and CoS2/OutofContract are serviced via the same queue. As such, they will have the same delay characteristics across the network. The difference is that in the event of severe congestion within CoS2, 'Out of Contract' class packets will be dropped first, allowing 'In Contract' CoS2 applications to be maintained.
- **CoS3 (In Contract) –** This class is indicated with DSCP Assured Forwarding 21 (AF21) and is intended for time sensitive, mission critical, bursty data applications.
- **CoS3 (Out of Contract) –** This class is indicated with DSCP Assured Forwarding 22 (AF22) and is intended for time sensitive, bursty data applications. CoS3/InContract and CoS3/OutofContract are serviced via the same queue. As such, they will have the same delay characteristics across the network. The difference is that in the event of severe congestion within CoS3, 'Out of Contract' class packets will be dropped first, allowing 'In Contract' CoS3 applications to be maintained.
- **CoS4 –** This class is indicated with DSCP default (default). It is also referred to as the best-effort class and is intended for all bulk data applications and non-time critical applications.

AT&T also stated that multiple traffic engineering techniques are applied to each application of CoS in order to help ensure an appropriate quality of service for each of the customer network applications. The management of latency and bandwidth in a customer network is accomplished by applying traffic shaping and bandwidth policing techniques, then assigning the traffic to an MPLS priority class on the customer router.

The techniques used to manage application traffic on the customer router, include:

1. Traffic is assigned to a class on the customer router. The parameters used to classify application traffic on the customer router, include:
   - Origin IP address
   - Destination IP address
   - Input interface
   - Port number
   - Application protocol
   - Classification/setting of IP precedence bits/marking
2. Traffic conditioning techniques include:
   - Classification/Setting of IP precedence bits/Marking
   - Traffic Policing and Traffic Shaping
   - Queuing Mechanisms
   - Congestion Control

The Customer Premises Equipment (CPE) uses this classification to differentiate the traffic and prioritize the applications at the CPE before transmission through the network. The objective is to optimize the access link utilization, as well as the service offered to the different types of applications. In the case of access congestion, high priority traffic takes precedence. For customers with customer-managed MPLS VPNs, AT&T makes available the ability to directly assign a Class of Service profile to traffic. AT&T states that this can be done without additional charges by using what AT&T refers to as BusinessDirect®.

Sprint

Sprint offers Class of Service as a standard feature for its MPLS solution. Customers do not pay additional fees or Monthly Recurring Charges (MRCs) and are not required to purchase Managed Network Services (MNS). Customers are provided flexibility in determining the number and size of their queues as Sprint does not place any restrictions. In addition, Sprint's network SLAs cover the entire port, not just certain queues. Sprint noted that most competitive MPLS solutions require additional fees, require MNS, and place restrictions on how many queues, the size of the queues, and what the SLA covers.

Service Level Agreements (SLAs)

There was a very wide variance in the level of information that The Carriers provided relative to the SLAs that they offer. Based on the information that was provided, the SLAs that The Carriers provide for their Layer 3 MPLS-based VPNs do not appear to show much innovation over what was available two or three years ago. In particular, it is still customary to have the SLAs be reactive in focus; i.e., the computation of an outage begins when the customer opens a trouble ticket. One of The Carriers even excludes from their availability target any network outage of less than a minute in duration. In addition, the level of compensation for violation of service level agreements remains quite modest.

In addition, The Carrier's SLA metrics are mostly still calculated as network-wide averages rather than for a specific customer's traffic or per site. As a result, it would be possible for a company's data center to receive poor service in spite of the fact that the network-wide SLA metrics remain within agreed bounds. In fact the propagation delay between two VPN sites could easily exceed the net average latency, so network-wide SLAs should not be used to set performance expectations. This means that enterprise IT organizations will probably need to gather performance metrics in addition to those offered by the carrier, in order to set appropriate performance expectations and ensure that SLAs are actually being met.

The reader is advised to query their vendors on the topic of customizing their SLAs. For example, AT&T stated that EVPN SLAs are specific to the customer's solution and can include end-to-end SLA's for specific site pairs.

Comparison and Observations

Table 1 provides a high level comparison of the BGP/MPLS VPN services offered by The Carriers. Note that Table 1 focuses on aspects of The Carrier's service offerings that were not already discussed. For example, Table 1 does not include a discussion of service classes.

Table 1 indicates that there are significant differences among The Carriers in terms of the:

- Protocols supported on the CE – PE interface
- Access options
- Value-added services (a.k.a. Advanced Services) that are accessible over the VPN
- VPN management functionality that is accessible via a web portal

Table 1: Comparison of BGP/MPLS VPNs

Vendor	AT&T	MCI	Sprint
Service Name/type	Enhanced VPN Service + VPN Transport Service RFC 2547bis	Private IP Service RFC 2547bis	MPLS VPN RFC 2547bis
Service Provider Manages the Router	Optional	Optional	Optional
CE-PE Route Exchange	Static, eBGP (standard) OSPF, EIGRP, RIP, RIPV2 (custom)	Static, eBGP, RIPV2, OSPF	Static, eBGP, OSPF, RIPV2 EiGRP
Access Options	Private Line, Frame Relay, ATM, Ethernet, DSL, WiFi, Cellular data, dial-up	Private Line, Frame Relay, ATM, Ethernet, DSL, Cable, Satellite	Private Line, DSL, Ethernet, Wireless (CDMA)
Coverage	50 states, 600 PoPs in US	Ubiquitous in the US	All cities in US
Intercarrier NNI Partnerships	2 PTTs in China, additional partnerships under consideration	Partnerships only in selected foreign countries; i.e., India	BellSouth, two with international carriers; future partnerships will likely be international
Pricing Structure - Recurring Costs	Access circuit + Port Speed + CoS Profile + Advanced Services	Access circuit + Port Speed + CoS Profile + Advanced Services	Access + Port Speed + Advanced Services; No extra charge for multicast or for CoS profile
Additional (Advanced) Services	Local, LD IPT/PSTN, IPSec Extranet integration, RAS integration Internet Access Network Based Firewalls Multicast 2H 2006	Local, LD IPT/PSTN, IP Multicast, Network Based Firewalls, IP Video Conferencing Bridging, IPSec Extranet integration (2006), RAS integration Internet Access	Local, LD IPT/PSTN, IP multicast (free) IPSec Extranet integration, RAS integration Internet Access
Web Portal Functionality	BusinessDirect Provisioning EBilling Reporting eTicketing Performance Stats Topology Map User self administration functionality	MCI Customer Center Provisioning eBilling Reporting eTicketing Performance Stats Topology Map	Multiple Names Utilization and Performance stats, ability to change class of service, some eBonding

Layer 2 MPLS-Based VPNs

As previously mentioned, none of The Carriers currently offers a Layer 2 MPLS-based VPN. However, each of The Carriers indicated that they are evaluating the possibility of leveraging their MPLS backbones to support Layer 2 VPN services.

In its simplest form the Layer 2 MPLS VPN emulates a Layer 2 point-to-point virtual circuit connection (a Pseudowire) between two CE routers or switches. This class of MPLS VPN is usually referred to as a Draft-Martini VPN or a Pseudowire Emulation (PWE) VPN. PWE can support a wide range of emulations, including Ethernet, Frame

Relay, ATM, HDLC, and PPP. PWE VPNs can be leveraged by the carriers to replace technology-specific switching elements (Frame Relay or ATM switches) in the core of the network. However, current versions of PWE do not scale very well because each pseudo-wire needs to be configured individually.

Another closely related Layer 2 MPLS VPN is the Virtual Private LAN Service (VPLS). The VPLS provides multipoint connectivity and enables extended LAN or LAN-to-LAN services. With VPLS, each customer's CE routers or switches appear to be attached to the same private LAN that comprises a customer-specific broadcast domain. In the Service Provider network, PE routers participating in a VPLS VPN are fully meshed in order to provide optimized site-to-site reachability without having to run a protocol such as spanning tree to avoid loops. Multicast traffic is treated as broadcast traffic, and hence is replicated across all PE ports that belong to a specific customer VPN instance. The full mesh and replication requirements can place limitations on the total number of VPLS PEs that can be deployed within a single VPLS domain. Once the total number of PEs is in the range of forty to sixty, some router vendors are recommending that Service Providers create a multi-tier hierarchy in order to scale VPLS services.

Within an MPLS/VPLS VPN, each CE is connected to the PE using Layer 2 attachment circuits (ACs). For Ethernet access to the VPLS, the attachment circuits can be specified by VLAN tags. A CE could also be attached to the PE via Frame Relay or ATM, in which case the attachment circuits would be designated by FR DLCIs or ATM VCs.

The PE maintains a Virtual Forwarder (VF) or Virtual Switch Instance (VSI) for each VPN. The VF performs service-specific processing and forwarding analogous to the VRF function for a Layer 3 MPLS VPN. The VF maps each AC to a Virtual Circuit (VC) and the VC to a tunnel LSP that has been traffic engineered for the traffic of a particular traffic class or forwarding equivalency class (FEC), e.g., real time traffic or general data traffic. IEEE 802.1p/Q VLAN IDs or tags in the customer's packets can be used to classify VC traffic into the proper tunnel LSPs.

A Pseudowire (PW) is a point-to-point path between two VFs in PEs belonging to a particular VPN. Autodiscovery and pseudowire signaling are used to find all the PEs participating in a VPN and to build a full mesh of PWs among the corresponding VFs. A PW is specified by the combination of a VC label and Tunnel label, which are both carried in the MPLS packet via label stacking. Autodiscovery is performed with BGP, while pseudowire signaling may be based on either LDP or BGP, both of which will be supported in the final IETF VPLS RFC.

Final Observations and Conclusions

Observations: State of the Marketplace

Each of The Carriers profiled in this report is in a period of great flux. For example:

- Verizon recently acquired MCI and is now branded Verizon Business. Verizon Business has stated that it will focus its MPLS offering on Private IP, the former MCI's flagship VPN service.
- SBC has recently completed its acquisition of AT&T.
- Sprint is merging with Nextel.

It is reasonable to expect that The Carriers will spend a lot of their energies in 2006 working on finalizing all of the organizational issues that are associated with the mergers and acquisitions that they announced in 2005. This means that a lot of their time will be spent trying to decide who works for whom, and where the real power resides in the new organizations. Given this, The Carriers' ability to implement innovation in their MPLS offerings will be slowed.

Conclusions

As was mentioned, none of The Carriers currently offers a Layer 2 MPLS-based VPN that is directly targeted at enterprise customers. It appears as if VPLS services from

The Carriers will not be introduced until 2006. Based on this and the fact that the VPLS RFC has yet to be finalized, it is likely that take two years or more for VPLS to become a mature, widely deployed VPN technology that can compete with Layer 3 MPLS-based VPNs.

In addition, VPLS may also prove to have scalability limits that prevent it from serving as a WAN backbone solution for large enterprises. If this turns out to be the case, VPLS may be relegated to the role of a metropolitan area network (MAN) technology.

For all of these reasons, Layer 2 MPLS-based VPN offerings may be too innovative for wide spread deployment in 2006.

Relative to the deployment of a Layer 3 MPLS-based VPN, section 2.0 of this document listed the primary factors that motivate companies to utilize these services. Those factors are:

- Combine multiple disparate networks onto a single network infrastructure
- Support any-to-any applications such as VoIP
- Enable efficient disaster recovery
- Prioritize applications in an easier fashion than is possible with Frame Relay or ATM
- Migrate off of legacy technologies such as Frame Relay and ATM in a seamless fashion
- Perform moves, adds and changes is easier
- Support for providing access to multiple data centers is easier

It is easy to conclude that Frame Relay and ATM are legacy technologies and that enterprise IT organizations must migrate off of these technologies at some point in time. It is also easy to conclude that at this time the successor technology for enterprise WANs is MPLS. The question becomes 'what is the right time for an enterprise IT organization to begin its migration from Frame Relay and ATM to MPLS?'

We recommend that IT organizations view the factors listed above as triggers. By triggers is meant that one or more of these factors should trigger an IT organization to begin its movement to implement MPLS. The most likely trigger is the need to support VoIP. A Layer-3 MPLS-based VPN would seem to be appropriate to support VoIP in part because all of The Carriers' current offerings have a class of service feature that is designed to carry voice traffic. That being said, given the relatively week SLAs offered by The Carriers, enterprise IT organizations need to assure themselves that the MPLS service offerings are robust enough to carry voice traffic.

Another reason that a Layer-3 MPLS-based VPN would seem to be appropriate is because VoIP tends to require any-to-any connectivity. In order to support this, the backbone needs to be highly meshed. A Layer 3 MPLS-based VPN outsources much of the complexity of managing full-mesh inter-site routing to the VPN service provider. It is also likely that with a Layer 3 MPLS-based VPN the cost of meshing will be reduced vs. what it would be with a Frame Relay or an ATM network because there is generally no additional cost for meshing or virtual circuits. In addition, multiple logical connections can share a single high-speed access line. However, any discussion of cost savings requires a detailed analysis of the company's traffic patterns combined with an analysis of the prices that they are paying for their current mix of services plus the price that they would be charged for a Layer 3 MPLS-based VPN.