Multiprotocol Label Switching for the Federal Government

Introduction

Service providers (SPs), enterprise, and federal customers are migrating from existing ATM, Frame Relay (FR), and Time Division Multiplex (TDM) infrastructures to an IP-based backbone. In the federal marketplace, this transition is under way for applications that service current e-government initiatives. In the Department of Defense (DoD) and the intelligence community (IC), there is a move toward the Netcentric Warfighter and next-generation (NG) IP applications. In each agency, the IP network is the key to providing these advanced services and applications.

As IP continues to evolve, so do the applications. Current IP backbones can no longer be designed to transport only IP packets. Instead, NG IP backbones must be capable of providing multiple IP services over a single physical infrastructure, using techniques such as quality of service (QoS) and security services. In addition, NG IP backbones should provide Layer 2/3 VPNs, IP multicast, IPv6, and granular traffic-engineering capabilities. Ultimately, these IP backbones should be scalable and flexible enough to support the mission-critical, time-sensitive applications that the federal government requires, and to meet new demands for applications, services, and bandwidth. Multiprotocol Label Switching (MPLS), when used on an IP backbone, provides the mechanism to offer rich IP service and transport capabilities to the router infrastructure.

This paper describes the architecture model being used to deploy NG IP networks and details the value of MPLS as an IP service and consolidation enabler for building these NG networks. It explains the services MPLS offers, provides deployment examples of each, and demonstrates how federal IT organizations and agencies can benefit from using MPLS as a service enabler. The goal of this paper is to provide positioning input to federal organizations and agencies as they migrate to their NG IP core infrastructures.

MPLS in the Federal Government

As federal customers embrace IP-based applications, more emphasis is being placed on the IP core infrastructures to support these new IP applications. MPLS is a key IP service-enabling option for offering and supporting multiple network applications and services. Various civilian, DoD, and IC agencies share network requirements that mimic small/large enterprise as well as SP networks. The following information describes ways that MPLS technology can benefit federal customers, such as customers who want to build their own MPLS networks or those who want to use MPLS L3 VPN technology as an offered transport service alternative to ATM/FR.

IP Services Overview

As IP networks continue to evolve, the applications and service demands required from the IP infrastructure evolve as well. SPs, federal IT agencies, and enterprises can no longer build their IP backbones to support only the fundamental transport of IP packets. The IP backbone must support multiple functions over the same physical infrastructure in order to reduce costs and ease manageability by consolidating the amount of devices in the network. This support reduces the operational complexities that exist when a single network is built per application, as shown in Figure 1.

Figure 1. Single Network per Application

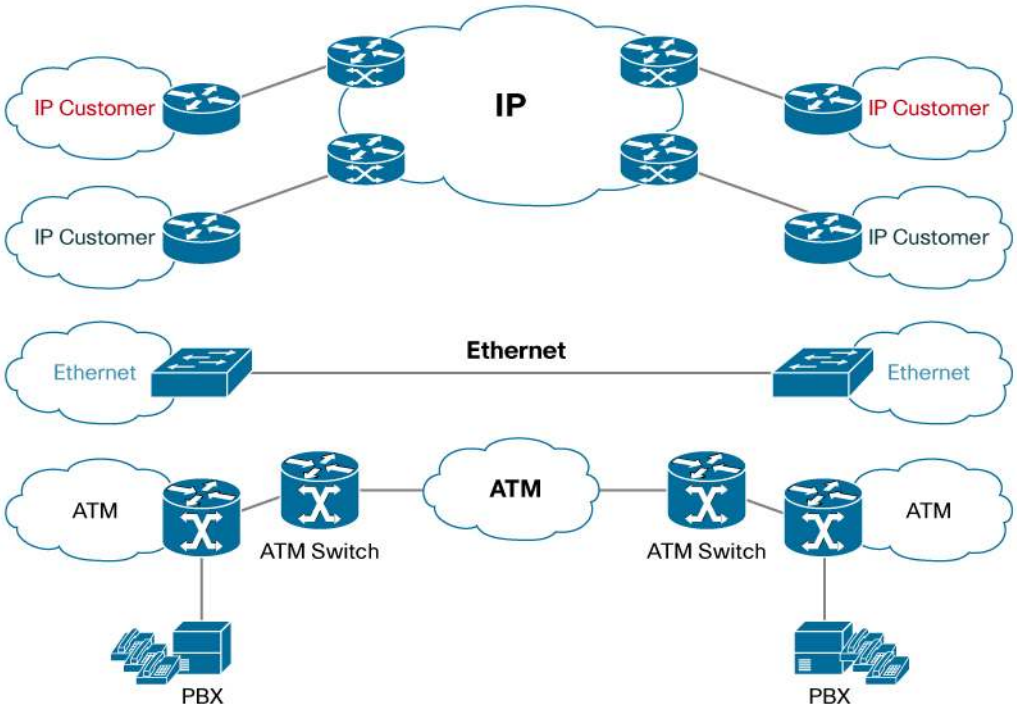

The IP core needs to support NG IP applications that are either already deployed or are being developed. These applications include voice, videoconferencing, streaming video, multicast applications, collaboration applications, mobility applications, and video surveillance. Each of these applications has certain network characteristics (e.g., delay, jitter, packet loss) that must be addressed so that each application can function as designed. In addition to applications, the network must support transport services capabilities, ideally over the same physical device/infrastructure, which will greatly reduce overall expenditures for operations and equipment. Network services can include Layer 2/3 VPN transport, multicast capabilities, support for differentiated services using QoS for both IPv4 and IPv6, as well as embedded security features and the capability to support centralized configuration, provisioning, and management.

As IT departments plan and architect for NG IP backbones, the capabilities of the network to support NG IP applications and offer the services listed above are crucial to building a consolidated multiservice IP backbone that is flexible and scalable well into the future. The next sections will provide high-level descriptions for the mechanism that will enable the IP backbone to support these NG IP-service capabilities. That mechanism is MPLS.

MPLS—The IP Service Enabler

In the early 1990s, Cisco® developed MPLS (originally called Tag Switching) to accelerate the performance of IP over an ATM infrastructure, but as time progressed, this technology evolved into a much more powerful consolidation and service enablement mechanism for IP backbones.

When enabled on an IP backbone, MPLS provides rich IP service and transport capabilities to the IP backbone for those federal customers either building their own MPLS backbone or planning to use a SP-offered Layer 3 VPN service.[1] MPLS can be used as a segmentation/virtualization method using Layer 3 VPNs in the campus/LAN/WAN. This method provides consolidation, cost

[1] The Layer 3 service referred to here is an IP VPN service. Whether the SP core is MPLS is irrelevant.

reduction, and a variety of transport features for IPv4, IPv6, and multicast within a VPN. MPLS also provides several Layer 2 transport capabilities through point-to-point Pseudowires and point-to-multipoint LAN services. In addition, MPLS can maximize the use of expensive WAN links and offer fast rerouting capabilities around link/node failures by enabling Traffic Engineering (TE) capabilities. Finally, combining MPLS TE with the QoS Differentiated Services (DiffServ) and Fast Reroute (FRR) capabilities allows for very granular QoS service offerings to end customers. Multiple service offerings delivered from a single router create a very effective form of migration and consolidation without the need of a network element per service.

Consolidation Using MPLS

In addition to the IP service enablement, MPLS offers tremendous consolidation capability with regard to device and cost reduction. Before MPLS, SPs required stovepipe networks per application (see Figure 1). For example, there was a separate infrastructure supporting TDM and PBX transport, a separate infrastructure for ATM and Frame Relay customers, and a separate infrastructure for IP transport. With each infrastructure came autonomous operation and management consoles as well as support personnel, all of which had huge operating cost implications. As MPLS evolved, so did the service enablement and consolidation capabilities, allowing the multiservice, multinetwork requirements to be consolidated onto a single IP infrastructure, as shown in Figure 2.

Figure 2. Consolidated Infrastructure Using MPLS

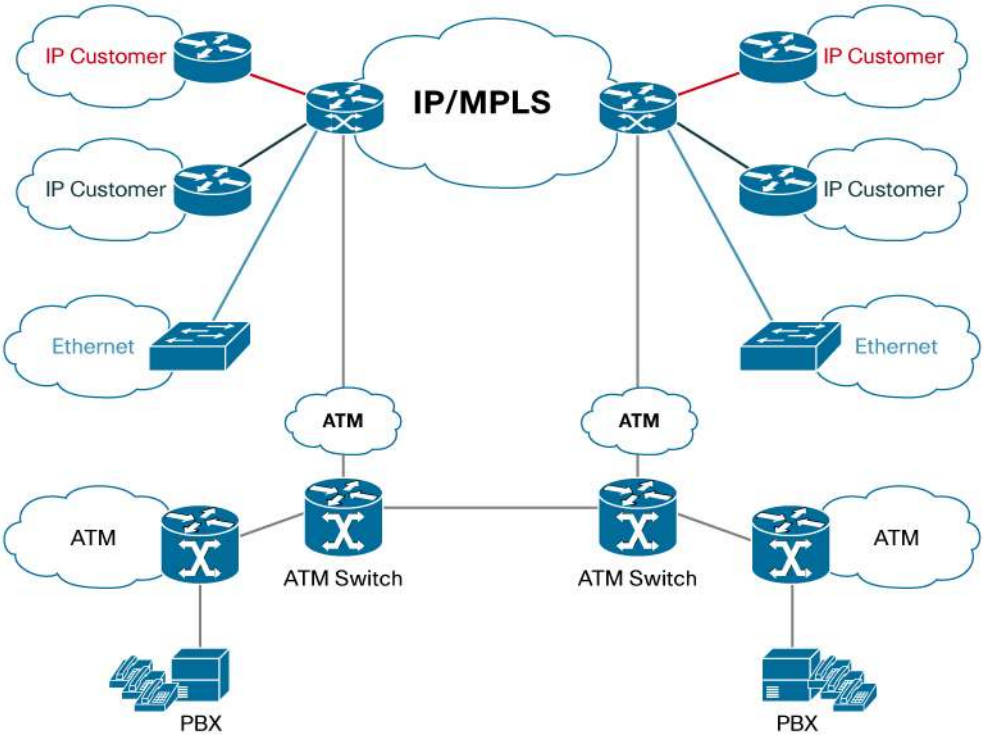

A single MPLS-enabled edge router can simultaneously support Layer 2/3 VPN and legacy encapsulation transport such as ATM/FR. This router can prioritize these different traffic types using QoS mechanisms, and provide SONET-like restoration times (~50 ms) around failed links and nodes in the MPLS core. This consolidation not only decreases the number of network devices to be managed, it also greatly reduces the operating expenses (OpEx) and capital

expenditures (CapEx) that are tied to the stovepipe networks, thus reducing the overall cost. This enormous cost savings has been the main factor driving MPLS into SP, enterprise, and federal IP backbones.

MPLS Service Options

In addition to consolidation features, service enablement also drives MPLS deployments. As discussed earlier, NG IP backbones must offer more than the transport of packets. These backbones need to provide a scalable set of IP services to accommodate NG IP applications being deployed. These backbones must also maintain the strict service levels (controlled latency, jitter, and packet loss) needed by these applications to maximize their potential for use. This section provides a description of these MPLS services and the benefits they can provide, not only to organizations that are SPs, but also to customers that are users of an MPLS service.

MPLS Services

Layer 3 VPN

Of all the MPLS services available, Layer 3 (L3) MPLS VPNs continue to be the most widely deployed by SPs across the globe. In this discussion, we will examine this service for:

- Federal customers building their own MPLS infrastructure and offering a L3 VPN service
- Federal customers using a SP-offered L3 VPN service.

Overview—MPLS Provider (Customer Building Their Own MPLS Core)

The MPLS L3 VPN model has several components, including the Provider (P), Provider Edge (PE), and Customer Edge (CE) routers. Normally, the MPLS core provider owns the P and PE routers, and the customer using the service owns the CE. (In a managed service, where the SP manages the entire solution, the SP also controls the CE.) The PE performs the primary functions as it is required to interface to the CE, exchange and hold routing information from the CE, and also peer with other PE routers in the MPLS core using Border Gateway Protocol (BGP). This arrangement is an extremely scalable solution, capable of supporting thousands of CE devices within a single MPLS infrastructure.

The keystone to the MPLS Layer 3 VPN is the Virtual Route Forwarding (VRF) component. A VRF instance can be thought of as a customer VPN or community of interest (COI) that provides a virtual IP routing table in a PE. VRFs can be unique to a specific customer and be capable of supporting thousands of routes being configured on hundreds of PEs throughout the network. It is important to note that the segmentation of each VRF maintains that no IP routes are shared between VRF, thus maintaining a separation mechanism equal to ATM/FR permanent virtual circuits. PE routers peer with each other over the MPLS core. Labels are used to distinguish the different VRFs on each PE. MPLS L3 VPNs can also support IP multicast, QoS, and IPv6 transport capabilities within a VRF. The MPLS-VPN deployments continue to increase in growth within many federal agencies, and this rapid pace will continue as the price model continues to drop.

Overview—L3 VPN Customer (Federal Customers Using an SP-Offered L3 VPN Service)

From the perspective of customers familiar with Layer 2 transport services (e.g., ATM/FR, T1/T3), the L3 VPN service offerings from SPs provide IP any-to-any connectivity. This eliminates the burden of IP scaling and peering issues, complex routing protocol issues, and having to constantly configure IP-to-Layer-2 mappings that overlay environments such as ATM and FR require.

In the traditional ATM/FR service (Figure 3), the SP does not participate in any of the customer's IP routing, thus creating an IP overlay network, and in a flat overlay transport creates an n-1 of IP peers.[2] The new IP L3 VPN service offering changes this paradigm to a peer model (Figure 4), as IP routing information is exchanged between the customer and the SP-owned PE router. The peer model puts the responsibility for distributing the IP routes and reachability for those routes from each CE on the SP network. This arrangement provides a major advantage as the L3 VPN SP provides any-to-any connectivity and segmentation, while eliminating the need for the customer to manage and maintain an entire IP core (including topology, routing protocol neighbors, core routers, and the operation of the entire infrastructure). The customer has only to manage the IP peering over the single CE to PE link.[3] This outsourcing of the core drives down the total cost of ownership to the end customer.

L3 VPNs continue to show rapid growth in all areas of the world. Thousands of federal government locations have shifted to the L3 VPN model versus the overlay model, greatly simplifying their operational model and reducing the complexities of managing a large IP core for agency transport.

Figure 3. Sample Overlay VPN Network

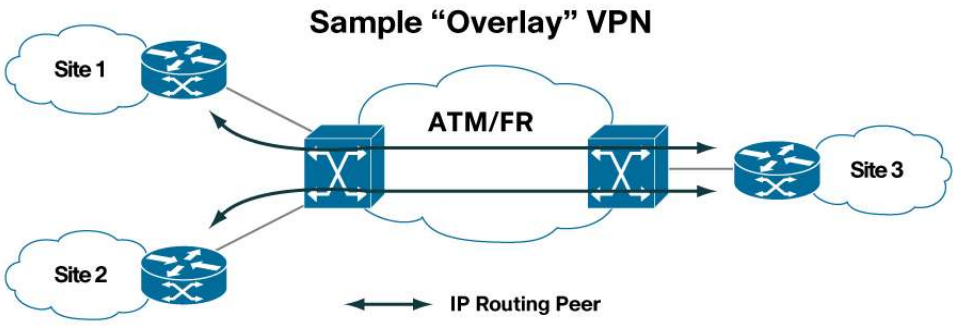

Figure 4. Sample Peer IP VPN Network

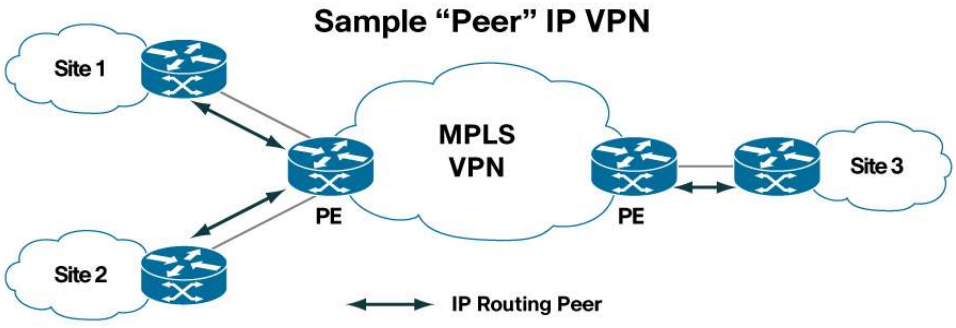

Layer 2 VPN and Pseudowire Transport

Overview

As the number of Layer 3 IP applications and service offerings continues to grow exponentially, the requirement for Layer 2 transport capabilities still exists in SP and federal networks. Many applications still require the transport of Layer 2 encapsulations over widespread locations, and customers do not want to build a separate infrastructure to support these requirements. For those

[2] For the ATM/FR PVC volume, (N*(N – 1) /2 applies.
[3] CE to PE protocols supported are RIPv2, OSPF, EIGRP, IS-IS, BGPv4, and static

requirements and others, MPLS-enabled IP backbones provide the mechanisms to offer and support Layer 2 transport capabilities over an IP-based infrastructure.

MPLS offers two primary solutions for Layer 2 services:[4]

- Virtual Private Wire Service (VPWS), which includes Ethernet, ATM/FR, Point-to-Point Protocol (PPP)/High-Level Data Link Control (HDLC) circuit transport
- Virtual Private LAN Service (VPLS), which offers point-to-multipoint LAN services over an MPLS-enabled IP core

VPWS, or Any Transport over MPLS (AToM), allows the transport of Layer 2 Pseudowire (PW) circuits (supporting the encapsulations listed above) over the MPLS core in a point-to-point fashion. In the federal government, AToM is being used to transport Ethernet VLANs over an IP backbone, specifically in a Metro environment. In the DoD and IC, the transport of ATM cells over an IP backbone is seen as a desirable transition solution when moving from ATM to IP backbones, and AToM is the mechanism that allows that capability.

AToM can use Metro Ethernet, which is a very cost-effective means of connectivity within a MAN. Plus, Ethernet is ubiquitous and easier to troubleshoot than ATM.

In addition to the point-to-point extension, MPLS offers a multipoint-to-multipoint VPN capability referred to as VPLS, which allows the MPLS backbone to emulate a virtual Ethernet switch function. This function interconnects multiple sites over great distances, as if each site was connected to a common LAN segment. This functionality can be useful for replacing LAN Emulation (LANE) networks, in the data center where legacy broadcast-only applications require Ethernet frame transport, as well as for broadcast applications that require any-to-any L2 transport that emulates a LAN segment over the WAN.

The benefit to using MPLS is that these Layer 2 features can be enabled on the same PE that provides Layer 3 services, thus eliminating the need for separate networks for each service offering and reducing the number of devices in the network backbone.

MPLS QoS

Overview

As MPLS continued to evolve and the number of deployments increased, it was critical that MPLS QoS capabilities remain superior so that customers using the VPN service could transport low-latency applications (i.e., voice and video over IP) with the proper priority. The MPLS backbone needed to have the capability to properly differentiate the traffic types in unison with IPv4 and IPv6. The backbone also needed to utilize the well-known QoS service markings that are used in non-MPLS LAN/WAN networks (e.g. CoS, ToS, and DSCP bits are each part of the IP packet header).

The MPLS label has the capabilities to support QoS markings through the use of the EXP, or Experimental bits, field. Customers who use an MPLS VPN service in which the SP provides a QoS offering must mark their traffic according to the specified markings the SP requests. The primary benefit is that customers using an MPLS service can have their enterprise QoS policy maintained even as the traffic traverses the MPLS core. This component supports next-generation voice and video applications.

[4] Each of these solutions is referenced out of the IETF PWE3 and L2VPN Working Groups (WG) respectively.

Traffic-Engineering Capabilities

MPLS Traffic Engineering (TE) was originally designed to overcome the inefficient use of links/bandwidth that is found in today's standard Interior Gateway Routing Protocols (IGRPs). For this discussion, there are three components that make up the TE architecture: fundamental TE, Fast Reroute (FRR), and DiffServ aware TE (DS-TE).

Overview

The basic function of MPLS TE[5] is to maximize link utilization within an MPLS core and avoid underutilization of other links in the core. This function is critical to reducing the link cost of the SP core networks and eliminates the issue of overutilizing or underutilizing links. MPLS TE components include the use of an Interior Gateway Protocol (IGP), such as Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (IS-IS), and each component contains extensions to properly flood the link information across the entire network and calculate a Constrained Based Shortest Path First (SPF) algorithm. Once the link attributes are flooded, MPLS TE requires the TE tunnel to be signaled across the network using Resource Reservation Protocol (RSVP)-TE, which calculates the parameters and distributes labels based on an algorithm.

Fast Reroute (FRR)

MPLS TE can also provide a FRR capability around failed links/nodes[6] in as little as 50 ms of the failure detection. This capability provides a very robust rerouting mechanism in the MPLS core. The basic function requires configuring preprovisioned tunnels that serve as backup to the primary tunnel carrying traffic. Figure 5 illustrates the FRR capability.

Figure 5. Fast Reroute

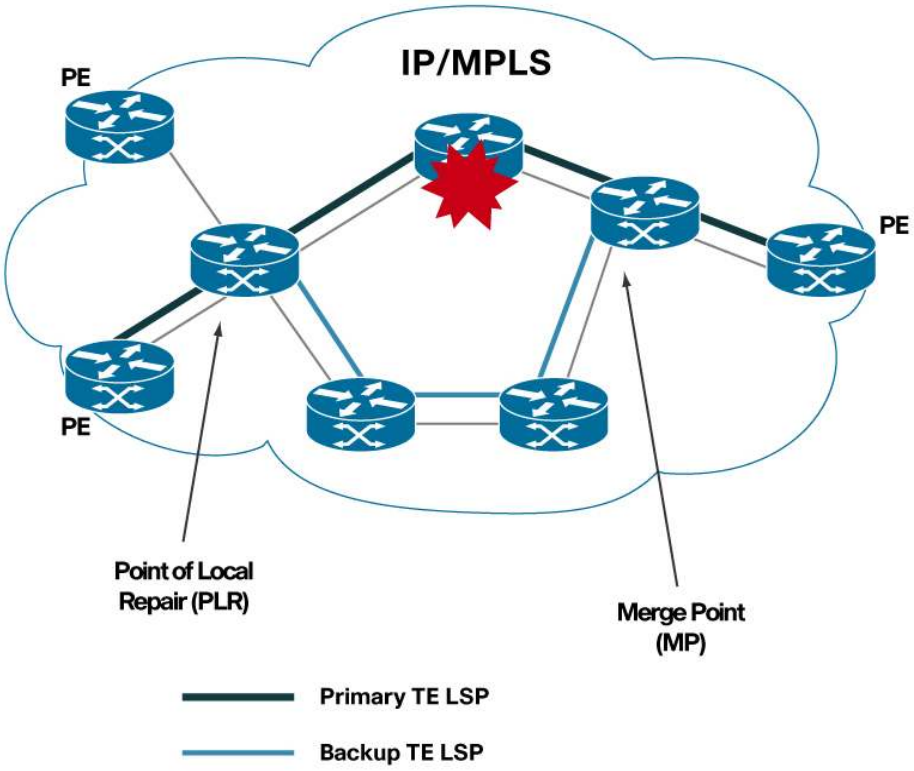

[5] Based on RFC 3290.
[6] MPLS TE also supports path protection mechanism as well, although it is not widely used.

Upon the failure of a link, node, or path in the network (a node in this example), the node where the backup tunnel is created—referred to as the point of local repair (PLR)—redirects traffic from the path that failed. Traffic flows to the pre-established backup tunnel that bypasses the node that failed, merging to a router downstream beyond the failure point (referred to as the merge point), all in a matter of milliseconds. Traffic will flow over the backup tunnel until the primary tunnel is restored, which is when traffic is redirected at the PLR.

DiffServ Aware TE

DiffServ aware TE is a mechanism, originally developed by Cisco, that allows a network operator to provide separate MPLS TE tunnels for a specific DiffServ class of service.

For example, in Figure 6, an MPLS TE tunnel can be configured to support the low-latency traffic class transporting voice, and the remaining best-effort traffic can use the TE tunnel. The low-latency tunnel can be configured to traverse the lowest latency path through the network core. This configuration assures the preferred path meets the packet-loss requirements as well as delay requirements for the voice application.

One use for DiffServ aware TE by a federal customer would be when a DoD or IC customer needs to avoid a high-latency satellite link for mission-critical applications that do not function well with a large amount of latency.

Figure 6. DiffServ Aware TE

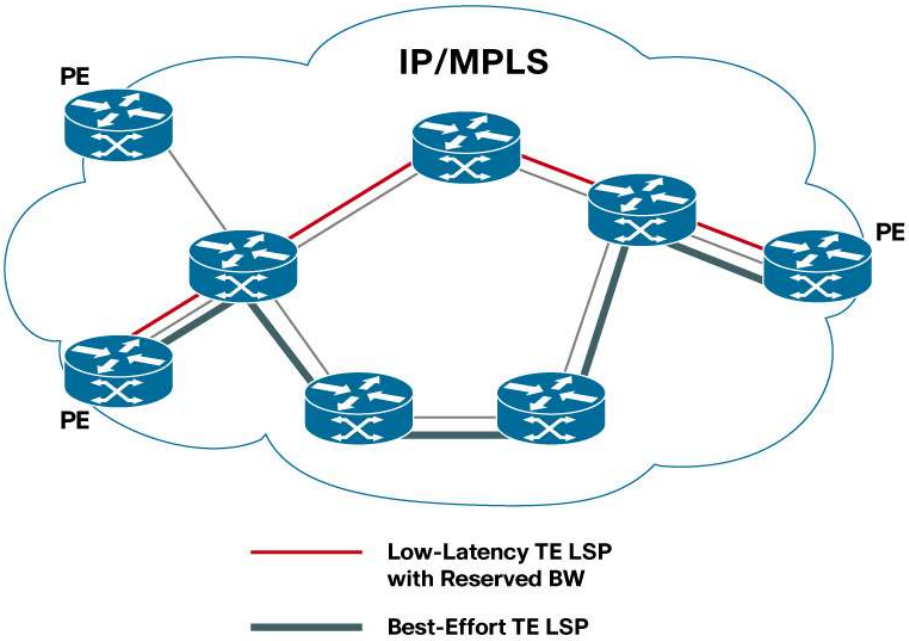

Assured Bandwidth Services

While not a specific feature, assured bandwidth services is the sum of the parts of multiple MPLS technologies previously discussed, specifically:

- MPLS TE
- MPLS QoS (DiffServ)
- MPLS DiffServ Aware TE (DS-TE)
- MPLS FRR.

Building upon the example in Figure 6, the combination of these QoS and protection features provide the assured bandwidth services capabilities. An example of where this combination is used is in the support of Virtual Leased Line (VLL) capability, shown in Figure 7.

Figure 7. Virtual Leased Line Solution

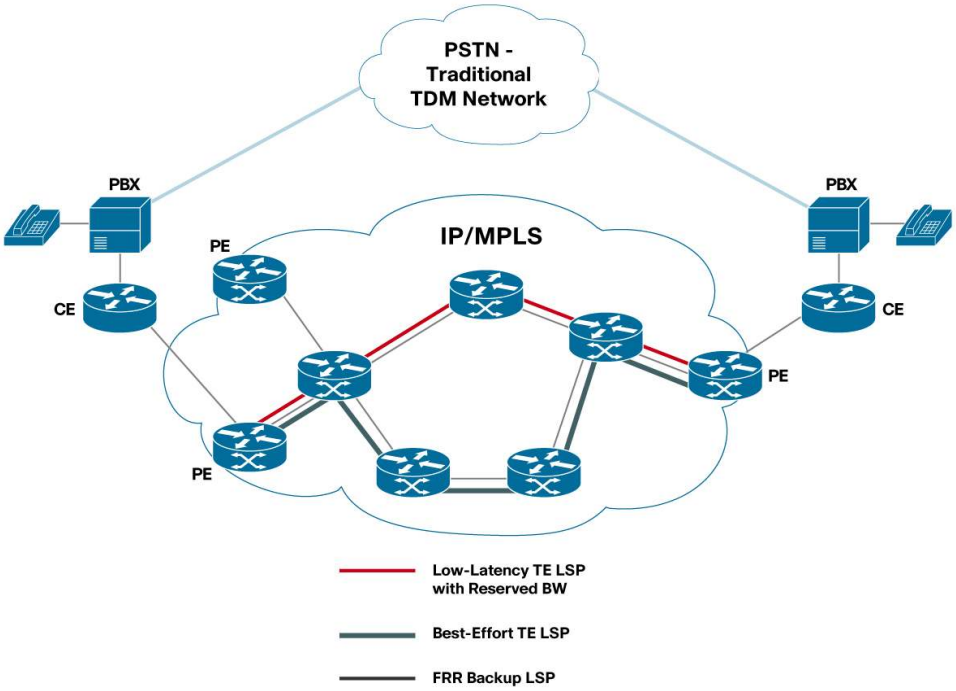

A VLL solution is responsible for carrying PBX trunk line traffic virtually across an MPLS backbone, and provides very stringent service-level agreements to customers that need low latency/jitter/loss requirements. In the case of VLL, the operator manually configures the low-latency TE tunnel path to traverse the link with the lowest amount of latency in the MPLS core (in this case, the top side of the core). The best-effort TE Label Switched Path is a lower quality path supporting best-effort traffic. In addition to manual configuration of the path, this low-latency TE tunnel would also be protected via the TE tunnel (FRR backup Label Switched Path) providing FRR node/link protection. In the event of a link or node failure along the low-latency TE Label Switched Path, FRR would provide less than 100-ms rerouting. This combination of technologies is what makes this assured bandwidth solution so powerful. Any federal customers requiring low-latency transport of specific applications with near-zero downtime requirements could benefit greatly from this type of service.

IPv6 Transport

As IP continues to be deployed to support next-generation applications and devices, IPv6 is beginning to gain traction throughout the world as a vehicle for technology innovation. IPv6 has gained much notice in the federal government, which has mandated that all agencies begin to migrate their network elements and operating systems toward an IPv6-ready state. For the federal IT managers, it is imperative that any agencies designing NG IP core networks be capable of supporting the transport of IPv6, and this also applies to MPLS backbones. MPLS technology has the capabilities of transporting IPv6 packets from PE to PE, in either a VPN or non-VPN configuration.

There are several advantages to having an MPLS-enabled core to transport IPv6 traffic, including:

- MPLS core devices (P routers) need no changes and no upgrades (they are unaware of IPv6).
- MPLS edge routers (PE routers) are the only devices needing upgrades to support IPv6 customer facing (or dual-stack if applicable).
- IPv6 transport can coexist and be supported simultaneously with already deployed features, as discussed above, in MPLS Services.
- 6PE/6VPE allows IPv6 to be deployed for support over the IP backbone, with minimal OpEx and CapEx. PE can support dual-stack (IPv4/v6) capabilities on the same physical interface and within the same VRF (for 6VPE).

Therefore, if an MPLS network is already in place, only minimal changes are required on the PE routers, thus minimizing overall core changes and impact to the operational network. SPs are also beginning to offer IPv6 as a transport service for basic transport, IPv6 internet access, and a dual-stack VPN service with IPv4.

Cisco IP/MPLS

Cisco IP/MPLS and supporting software technologies feature a robust and flexible networking technology based on the powerful Cisco Intelligent Information Network (IIN) solutions. Cisco IP/MPLS delivers unparalleled control, data, and management plane scalability; continuous system operation; unprecedented service flexibility; end-to-end intelligent management and security; and global 24-hour support from Cisco and Cisco-certified partners. Cisco IP/MPLS is the IP next-generation network and IIN convergence platform that enables service providers and enterprises to build intelligent networks through its resilient, integrated, and adaptive capabilities. Cisco IP/MPLS enables customers to optimize their business through network convergence and simplification, grow their business through innovative and differentiated services, and protect their business for operational efficiency and profitability.

Why Cisco?

The undisputed technology of choice and market leader, Cisco IP/MPLS runs in:

- More than 95 percent of service provider customer networks—close to 300 worldwide
- The first do-it-yourself global enterprise customers networks

As the industry's premier innovator, Cisco has advanced IP/MPLS technology with new features and functionality, including Tag Switching and, most recently, AToM. With the world's broadest platform support and largest installed base, Cisco continues to be the industry leader in IP/MPLS.

Summary

Next-generation IP networks must support an abundance of services beyond IP packet transport. MPLS is vital to service enabling the IP core and providing a level of consolidation. It is critical that the network be capable of supporting next-generation IP applications, many of which require low-latency transport with minimal jitter and packet loss. MPLS can offer a variation of segmentation at Layer 2 and Layer 3, QoS, IPv6 transport, and resiliency at the IP layer that is equal to SONET restoration. In addition, the combination of multiple MPLS mechanisms can create a zero downtime model that protects applications with almost no packet loss. As IP core networks are deployed, it is critical that the design allows multiple service offerings in a scalable manner without building parallel networks per application. This design reduces the overall OpEx and CapEx of supporting the network and simplifies the management process.

Cisco provides the industry's only complete IP/MPLS solution and addresses both service provider and enterprise customers' end-to-end needs—from managing bandwidth and traffic in the core to VPN services in the edge. Subsequent Cisco IP/MPLS White Papers will detail how several topics discussed in this paper apply to the federal government's requirements.

Acronyms	
BGP	Border Gateway Protocol
CapEx	Capital Expenditures
CE	Customer Edge
COI	Community of Interest
CoS	Cost of Service
DiffServ	Differentiated Services
DoD	Department of Defense
DS	DiffServ
DiffServ	Differentiated Services
DoD	Department of Defense
DS	DiffServ
DSCP	Differentiated Services Code Point
EIGRP	Enhanced Interior Gateway Routing Protocol
EXP	Experimental Bits (Field)
FR	Frame Relay
FRR	Fast Reroute
HDLC	High-Level Data Link Control
IC	Intelligence Community
IGP	Interior Gateway Protocol
IS-IS	Intermediate System to Intermediate System
L2 VPN	Layer 2 Virtual Private Network
LANE	Local Area Network Emulation
LDP	Label Distribution Protocol
MPLS	Multiprotocol Label Switching
NG	Next Generation
OpEx	Operating Expenses
OSPF	Open Shortest Path First
PBX	Private Branch Exchange
PE	Provider Edge

PLR	Point of Local Repair
PPP	Point-to-Point Protocol
PW	Pseudowire
PWE3	Pseudowire Emulation Edge to Edge
QoS	Quality of Service
RFC	Remote Function Call
RSVP	Resource Reservation Protocol
SLA	Service Level Agreement
SP	Service Provider
SPF	Shortest Path First
TDM	Time Division Multiplex
TE	Traffic Engineering
ToS	Type of Service
VLL	Virtual Leased Line
VPLS	Virtual Private LAN Service
VPWS	Virtual Private Wire Service
VRF	Virtual Route Forwarding

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Printed in USA

C11-390659-00 03/07