

MPLS in the Enterprise

Introduction

Multi-Protocol Label Switching (MPLS) as a technology has been around for over a decade and has extensively been used in several service provider networks world-wide. Over the last few years, the standardization of applications such as VPN technologies over MPLS has opened the door for this technology to be used in an enterprise network.

This paper will explore the use of MPLS in an enterprise network and its associated benefits.

Overview of MPLS

MPLS was developed in the late 90s as a means for rapid switching of packets in an IP network. In contrast to the connectionless model used in IP networks, MPLS establishes a Label Switched Path (LSP) into which packets are encapsulated by an edge MPLS router after assigning a 20-bit label. Subsequent nodes (called Label Switching Routers) forward the packet by inspecting and swapping the label.

Packets that are forwarded in the same manner (e.g. those with the same longest prefix match) are considered to belong to the same "Forwarding Equivalence Class" (FEC) and are therefore assigned the same label. Figure 1 depicts how a packet is forwarded within an MPLS network.

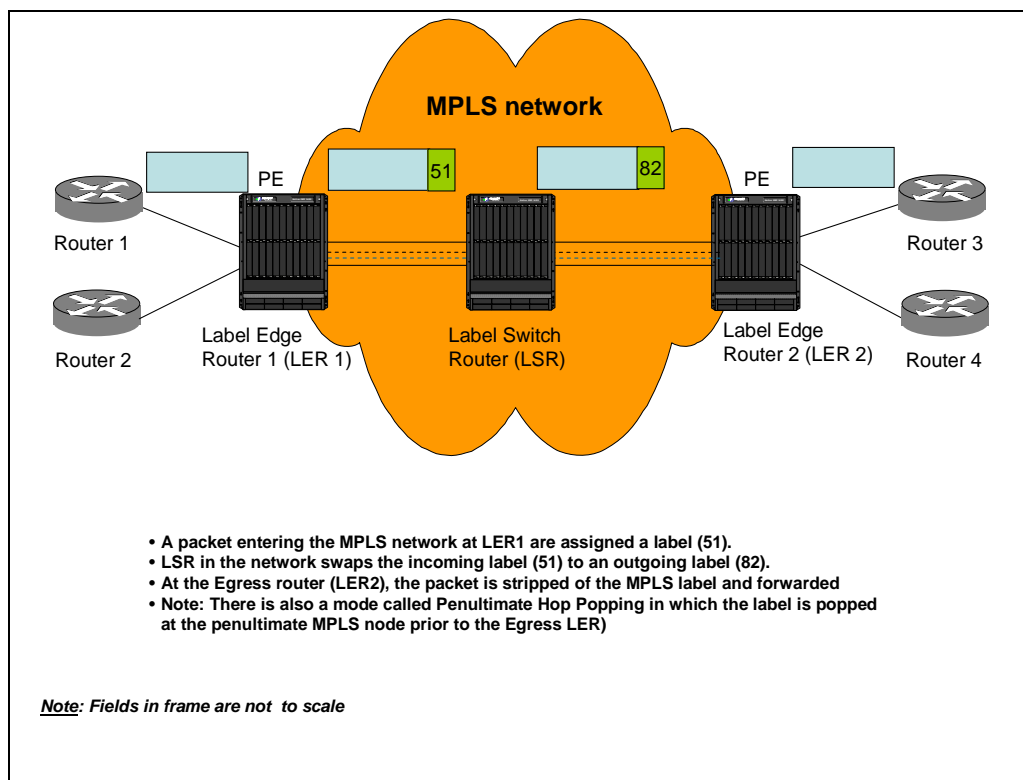


Figure 1: Forwarding of a packet in an MPLS Network

To understand the operation of an MPLS router, it is useful to understand 3 building blocks:

- Assign the packet to a FEC. In order to do this, network reachability information is required and is propagated using a conventional routing protocol such as OSPF.

- Assign a label to packets that map to the same FEC and propagate it to other nodes in the network. Label propagation is done using an MPLS signaling protocol such as LDP (Label Distribution Protocol) or RSVP (Resource Reservation Protocol).
- To get the best benefits of MPLS, traffic engineering extensions to the routing protocol are used to propagate information on available bandwidth and administrative constraints to all nodes in the MPLS network.

The first 2 steps are mandatory in an MPLS network. The use of traffic engineering is optional.

Over time, MPLS technology has evolved to support tunneling mechanisms, including the ability to establish Virtual Private Networks (VPNs). Because forwarding is done based on labels within an MPLS network, the use of multiple nested labels provides a method to tunnel traffic from multiple instances within a common outer label.

Benefits of MPLS

There are several benefits that MPLS offers:

- Its ability to efficiently label switch packets allows multiple services to be easily supported on a converged infrastructure. Unlike ATM or Frame Relay, MPLS is tightly integrated with IP—the assignment of the MPLS label is based on the FEC as described above. This leads to a better integration with IP networks.
- MPLS' connection-oriented model and rich support for traffic engineering, allows specific preferred paths to be taken for certain types of traffic, resulting in very high QoS being delivered.
- The ability to offer sub-50 millisecond protection in the event of failure using a method called Fast Re-Route allows rapid convergence to be achieved, allowing quick re-routing around a failure. This is particularly useful when carrying VoIP or mission-critical traffic in an enterprise network.
- With the standardization of VPN technologies over MPLS, large Layer 2 domains can be created without the traditional disadvantages of a Layer 2 design.
- The technology also allows Layer 3 VPNs to be created, with multiple administrative domains based on Layer 3 MPLS VPNs.

The last 2 capabilities are of particular interest because multiple types of VPN technologies can be delivered concurrently over an MPLS network. By combining this with an Ethernet-based infrastructure, a high-performance network can be easily created and managed.

Layer 2 Service without the Drawbacks of a Layer 2 Network

Enterprise networks have traditionally been built with a switched architecture at the edge interconnected using a core routing infrastructure. Nevertheless, there are situations when it is necessary to have a large Layer 2 span.

For example, it is common to have different departments or administrative groups segregated by VLANs within an enterprise. When users who belong to these departments are not in close geographical proximity, it may require VLANs to span across a large section of the enterprise network and in some cases, even across the core of a network. IT administrators often despise running large Layer 2 sections across the core of a network and therefore typically set aside separate switches to provide such connectivity. Figure 2(a) depicts an example of such a network.

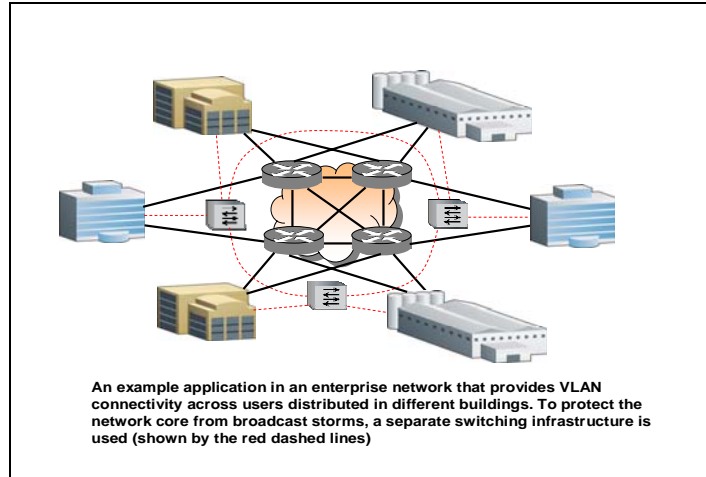


Figure 2 (a): Separate network used for VLAN connectivity among users in different buildings

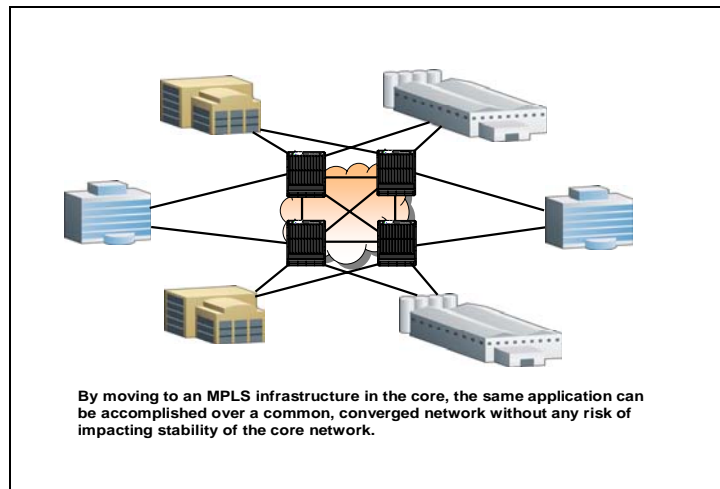


Figure 2 (b): Use of an MPLS network provides resiliency and a way to interconnect users at Layer 2 without any risk of impacting stability of the network core

A second example is when multiple branch offices of a large corporation need to be interconnected. Historically, geographically separated networks such as branch offices of a large corporation have been connected by leased T1 lines, Frame Relay or ATM connections. With the progress in Carrier-Grade Ethernet, such connectivity can well be accomplished using an Ethernet service.

The standardization of VPN technologies over MPLS allows such applications to be easily delivered using MPLS. Three types of VPNs can be built over an MPLS network:

- A point-to-point Ethernet service provides connectivity between any 2 nodes in an Ethernet network and is commonly called Virtual Leased Line (VLL) or Virtual Private Wire Service (VPWS)
- A multipoint Ethernet service offers connectivity among multiple sites using a method called Virtual Private LAN Service (VPLS). Thus, a VPLS emulates the behavior of a traditional IEEE 802.1D bridge over an MPLS network.
- A Layer 3 VPN service that allows multiple virtual domains to be set up using a method called BGP/MPLS VPNs. This technology uses a combination of BGP and MPLS for creating multiple Layer 3 VPNs.

Of these, VPLS and VLL are particularly interesting to enterprises because they provide effective alternatives to IT administrators averse to deploying loop-mitigating protocols such as RSTP or MSTP across large Layer 2 domains. These mechanisms work by establishing a “pseudo-wire” between participating routers and using a second label for carrying traffic within the VPLS or VLL. For VPLS, a full-mesh of pseudo-wires is established between participating MPLS routers to allow the Ethernet frame to be forwarded to the right destination.

When processing broadcast or unknown unicast frames, VPLS uses a split-horizon method to prevent loops within the network. There is therefore no need for a loop-mitigating protocol (such as RSTP or MSTP) within the domain of the MPLS network.¹

Applicability of MPLS in an Enterprise Network

There are several examples where MPLS is applicable in an enterprise network:

- Enterprises seeking to provide connectivity among several sites and who either own or have the ability to affordably lease fiber to connect these sites (e.g. a school district, utility provider, City Government, an enterprise that own multi-mode fiber connecting their different buildings)
- An enterprise network that has many VLANs spanning a large portion of their network
- Enterprises looking to migrate from FDDI/ATM networks installed in the 1990s
- An enterprise seeking to install a converged network

Administration of an MPLS Network

Administration of an MPLS network is surprisingly simple, particularly in an enterprise network. Because an MPLS network is built on top of an IP network, administration is very intuitive and there are well established tools for managing and monitoring an MPLS network.

A Deployment Example

Washington School District (WSD), a school district in Eastern Missouri, US, recently implemented a successful transition to an MPLS network. The District initially had an assortment of frame relay circuits to provide connectivity across the school district in addition to several switches and routers in their network. In addition to fast convergence and the ability to create TLS (Transparent LAN Service) tunnels across the MPLS network, 2 other factors were crucial in their choice of an MPLS network:

- Addition of new TLS service instances needs to be done only at the edge of the network without requiring any changes in the network core. In contrast, running VLANs across the network core would require making changes to the network core for every new service added.
- In the past, WSD had broadcast storms that affected the entire network by impacting the CPU load on the original L2/L3 backbone switches. With the use of VPLS, traffic is transparently sent to the members of the VPLS, without affecting the core of the network.

Dale Loesing, Network Administrator at Washington School District says, ““We moved to an MPLS backbone utilizing FRR to achieve faster (50 ms) convergence for our VoIP traffic across our entire school district. What we got was a simpler, easier to manage network that I don’t have to touch as much anymore. Using VLL and VPLS tunnels has eased the management overhead and improved our backbone stability by eliminating configuration changes each time I need to extend a VLAN across the network.”

A large multi-site hospital in Canada similarly migrated to an MPLS network recently with enormous success.

¹ For a more detailed explanation of VPLS and VLL, refer to the white paper “Offering Scalable Layer 2 Services With VPLS and VLL” at <http://www.foundrynet.com/pdf/an-offering-scalable-l2-services-vpls-vll.pdf>

Conclusions

Enterprises seeking to build a large converged network that supports multiple services could benefit by considering the use of MPLS within their network. It is a proven technology that has been implemented in networks both large and small on a global scale.

Author: Ananda Rajagopal
Document version 1.0

Foundry Networks, Inc.
Headquarters
4980 Great America Parkway
Santa Clara, CA 95054-1200
U.S. and Canada Toll-free: (888) TURBOLAN
Direct telephone: +1 408.207.1700
Email: info@foundrynet.com
Web: <http://www.foundrynet.com>

Foundry Networks, AccessIron, BigIron, EdgeIron, FastIron, IronPoint, IronView, IronWare, JetCore, NetIron, ServerIron, Terathon, TurboIron, and the "Iron" family of marks are trademarks or registered trademarks of Foundry Networks, Inc. in United States and other countries. All other trademarks are the properties of their respective owners.

Although Foundry has attempted to provide accurate information in these materials, Foundry assumes no legal responsibility for the accuracy or completeness of the information. More specific information is available on request from Foundry. Please note that Foundry's product information does not constitute or contain any guarantee, warranty or legally binding representation, unless expressly identified as such in a duly signed writing.

©2007-2008 Foundry Networks, Inc. All Rights Reserved