Layer 3 Switching Using MPLS

White Paper

Important Notice

Trademark Acknowledgements

Harris & Jeffries, Inc.

Introduction

Multiprotocol Label Switching (MPLS) represents the next level of standards-based evolution in combining layer 2 (data link layer) switching technologies with layer 3 (network layer) routing technologies. The primary objective of this standardization process is to create a flexible networking fabric that provides increased performance and scalability. This includes traffic engineering capabilities which provide, for example, aspects of Quality of Service (QoS) /Class of Service (CoS) and facilitate the use of Virtual Private Networks (VPNs).

MPLS is designed to work with a wide variety of transport mechanisms; however, the initial implementations will focus on leveraging ATM and frame relay which are already deployed in large-scale service provider networks.

This white paper highlights the challenges faced by service providers in transitioning to IP-based networks, the current capabilities of IP, the current methods used to build large-scale IP networks, and the capabilities envisioned for the deployment of MPLS that will address the requirements of the service provider community.

This guide, the first in a series of MPLS technology guides, provides an overview of MPLS and its basic concepts. Future guides will lend insight into the current state of the MPLS standardization process, technical details describing MPLS operation, and an overview of how Harris & Jeffries is assisting equipment manufacturers to meet the evolving market demands.

MPLS Background

MPLS is being defined by the Internet Engineering Task Force (IETF) as a standards-based approach to applying label switching technology to large-scale networks. The MPLS Working Group was established in early 1997 and has since defined a large set of working documents that are currently being massaged into standards. Prior to the formation of the MPLS Working Group, a number of vendors had announced and/or built a proprietary version of a label switching implementation. It is this widespread interest in label switching technology that initiated the formation of the MPLS Working Group.

The IETF is defining MPLS in response to numerous interrelated problems that need immediate attention. These problems include, scaling IP networks to meet the growing demands of Internet traffic, enabling differentiated levels of IP-based services to be provisioned, merging disparate traffic types onto a single IP network, and improving operational efficiency in a competitive environment. It is important to note that many service providers are active in the MPLS Working Group. This ensures that the capabilities of MPLS will have a direct correlation to the aforementioned problems.

Many of the issues MPLS is seeking to address have already been recognized by a number of equipment manufacturers. In fact, many of these vendors have already developed proprietary solutions that address these problems. These include, IP Switching from Ipsilon (Nokia), Tag Switching from Cisco Systems, Aggregate Route-based IP Switching (ARIS) from IBM, Cell Switch Router (CSR) from Toshiba and IP Navigator from Cascade Communications (Ascend Communications). The MPLS documents draw on these existing implementations in an effort to produce a widely applicable standard.

Although MPLS, as its name suggests, can conceptually support multiple protocols, the initial work is focused on the integration of IPv4 with ATM and frame relay.

MPLS Overview

ATM was once envisioned to be a ubiquitous technology. Many people thought that ATM would extend from the desktop, through the core of the network and terminate at another desktop; in other words, a homogeneous ATM network. Today, few people continue to embrace that vision. Instead, homogeneous networks utilizing IP as the transport mechanism are being built and are likely to become commonplace over time. However, economics will play an important role in determining the adoption rate of these next generation IP networks. Upgrading entire networks by swapping out hundreds or even thousands of pieces of equipment is cost prohibitive. Therefore, many service providers will continue to maintain ATM in their networks for the foreseeable future. Consequently, next generation networks will be built using new technologies that leverage existing technologies such as IP and ATM, which are already deployed in existing networks.

Label switching technology is a result of the desire to combine the benefits of switching technologies that live in the core of the network with the benefits of IP routing technologies that live at the edge of the network. A hybrid network utilizing both of these technologies creates a problem best described as "how to make IP and ATM interoperate". The IETF and the ATM Forum initially took up this challenge and defined standards such as IP over ATM (RFC 1577/2225) and Multiprotocol over ATM (MPOA) that enable IP to work over an ATM network. However, MPLS is tackling a different problem best described as "how to integrate the best attributes of traditional layer 2 and layer 3 technologies".

As previously mentioned, label switching seeks to combine the best attributes of layer 2 switching technologies, as embodied within the ATM and frame relay world, with the best attributes of the layer 3 routing technologies as embodied in the IP world. MPLS, as the standards-based approach to label switching, specifically defines a set of protocols and procedures that enable the fast switching capabilities of ATM and frame relay to be utilized by IP networks.

The key concept in MPLS is identifying and marking IP packets with labels and forwarding them to a modified switch or router, which then uses the labels to switch the packets through the network. The labels are created and assigned to IP packets based upon the information gathered from existing IP routing protocols.

As with any technology, MPLS has its own set of terminology and acronyms. Some of the key terminology and concepts (figure 1) should be defined before proceeding.

- *Datagram* is a generic term that refers to a unit of information being sent through a network.

- *Label edge router (LER)* is a device that sits at the edge of an MPLS domain and is capable of utilizing the routing information to assign labels to datagrams and then forward them into an MPLS domain.

- *Label switched path (LSP)* is a specific path that a datagram travels through a network based upon the labels that are assigned to that datagram.

- *Label switching* is the term used to describe the generic technology that combines layer 2 and layer 3 technologies.

- *Label switching router (LSR)* is a device that typically resides somewhere in the middle of a network and is capable of forwarding datagrams based upon a label. In many cases, especially early versions of MPLS-based networks, a LSR will typically be a modified ATM switch that forwards datagrams based upon a label residing in the VPI/VCI field.

- *MPLS* is the term used to describe the specific work going on in the IETF to standardize label switching.

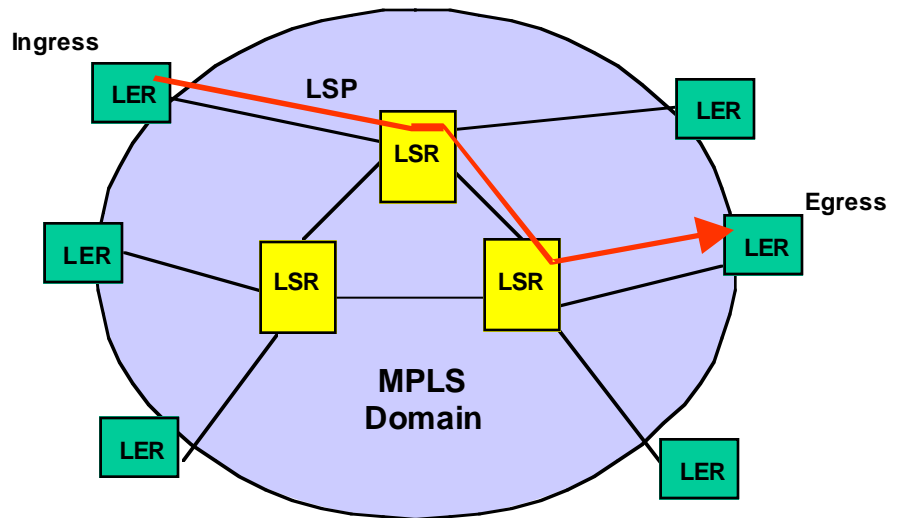- *MPLS domain* is a portion of a network that contains devices that understand MPLS.

Figure 1

Service Provider Challenges

Deregulation of the telecommunications industry, particularly in North America, has led to a proliferation of alternative service providers who compete in market niches, both with each other and with the more traditional carriers.

This highly competitive business environment is driving service providers to closely examine their network cost structures, including both infrastructure costs and operational costs. Providing a homogeneous network transport mechanism is viewed as the most cost efficient and flexible way to address these issues. Increased competition has also led service providers to seek innovative ways to expand their service portfolios so that they can differentiate themselves in the marketplace and target specific market segments.

A wide variety of services are available from the service provider community today. These services have evolved over many years with specialized networks deployed to meet specific requirements. Enterprises are no longer content to deal with multiple network connections and are demanding hybrid service offerings with a single flexible access, and in the future, a uniform access protocol, such as IP.

Although traditional telecommunications services and more recent data services, such as Intranets/Internet access, can be converged on a single network, enterprises are increasingly looking for more advanced and specialized services tailored to their specific needs. This will enable enterprises to dynamically adjust their network requirements based upon factors such as traffic loading per application (bandwidth allocation) and application performance (QoS/CoS).

Current IP networks are a long way from meeting the requirements of service providers and their customers. The capabilities embodied in MPLS are designed to meet the unified transport requirements of large-scale IP networks and build on the existing concepts of IP networking.

Public Networks Today

Growth of IP Networking

The origin of IP networking dates back to the 1970s in the academic community. At the time, the primary motivation for packet communications was robust information exchange. Delivery of packets with specific bandwidth or delay requirements was not a factor. Even though IP embodies the concept of differentiated packet transport via the type of service (TOS) categories, this capability has never been utilized with any consistency. Consequently, IP networks and the Internet have been built using a connectionless packet switching model that doesn't provide any quality of service differentiation.

The growth of the Internet and the emergence of the World Wide Web (WWW) have propelled the Internet Protocol (IP) to the forefront of the communications world. Both public networks, primarily the Internet, and private networks (corporate Intranets) use IP as the foundation for their data networking needs.   Many service providers are now looking to build or consolidate their network operations on a single infrastructure with IP-based networks being the network of choice. However, the transport requirements of voice, data and multimedia applications differ widely and the architecture of TCP/IP and the IP protocol itself need to evolve to meet these diverse demands.

The Role of Routing

IP packets contain a header with sufficient information that enables them to be forwarded through a network. Packet forwarding has traditionally been based upon datagram routing. The datagram routing technique used in IP networks is a destination-based routing paradigm. This means that an IP packet is routed through the network based upon the destination address contained within the packet. The forwarding mechanism utilized by IP networks is "hop-by-hop" routing, which means that every packet entering a router is examined and a decision is made as to where to send the packet (i.e., what is the packet's "next hop"). In this manner, a packet is routed through a network from its source to its destination. Since the packets are individually routed through a network and don't follow a predetermined path, the network is said to be connectionless.

Routers exchange information by establishing an adjacency (i.e., conversation) with every directly connected peer. The example in figure 2 shows that Router A has three directly connected peers (B, C, and D) meaning that Router A requires three adjacencies. Similarly, Router C has two directly connected peers (A and D) so Router C only requires two adjacencies.
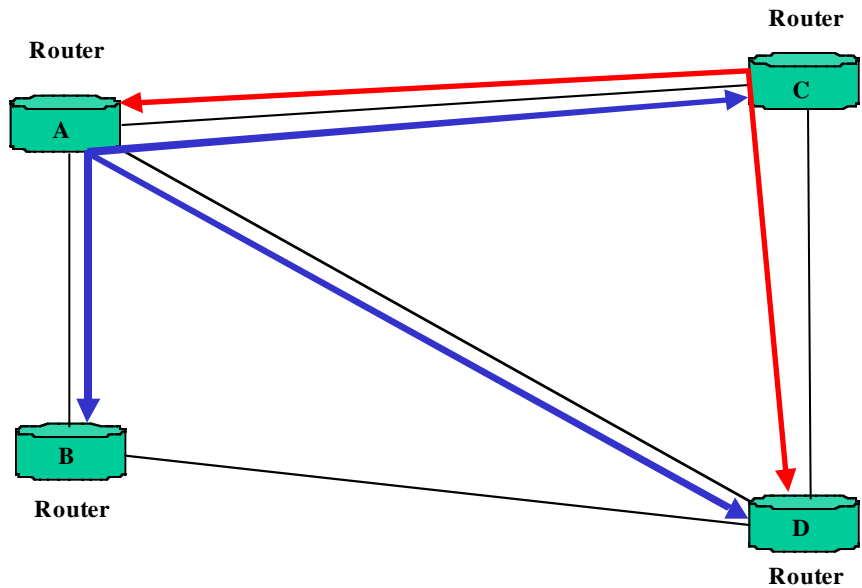
Figure 2

In order to properly route a packet, a router must be able to determine the next hop for a packet. Routing protocols, such as Open Shortest Path First (OSPF), enable each router to learn the topology of the network. Using the information provided by the routing protocols, the routers build forwarding tables that identify the appropriate next hop for all known IP destination addresses. Note: routers typically store IP prefixes rather than complete IP addresses in their forwarding tables; however, the details of IP routing are outside the scope of this paper. The construction of the forwarding tables and the forwarding mechanisms, although closely related, are separate functions (figure 3). Once implemented purely in software, the forwarding functions are now being implemented in hardware with the routing/control functions remaining in software.
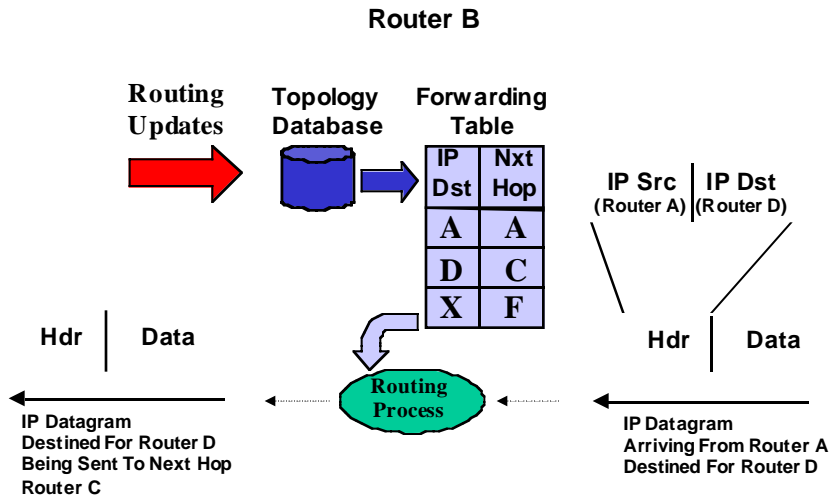
Router B

Routing Updates | **Topology Database** | **Forwarding Table**

IP Dst	Nxt Hop
A	A
D	C
X	F

IP Src (Router A) | **IP Dst** (Router D)

Hdr | **Data**

Hdr | **Data**

Routing Process

IP Datagram
Destined For Router D
Being Sent To Next Hop
Router C

IP Datagram
Arriving From Router A
Destined For Router D
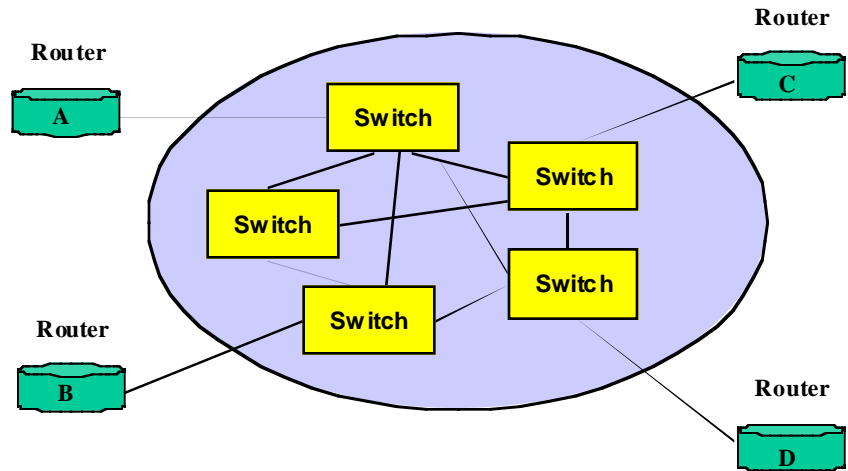
Figure 3

The Role of Switching

As service providers and enterprises developed large IP networks, they soon realized that building router-based networks created a number of issues. These issues were largely related to the software forwarding component of IP routers, the expense of high speed router components themselves and the difficulty in predicting performance in a large meshed network based upon traditional routing concepts.

Switching technologies based upon ATM and frame relay utilize a much different forwarding algorithm that is essentially a label-swapping algorithm. Since this forwarding algorithm is so simplistic, it is typically done in hardware, yielding a better price/performance advantage when compared to traditional IP routing. ATM and frame relay are connection-oriented technologies, meaning that traffic is only sent between two endpoints once a connection (i.e., a pre-determined path) has been established. Since traffic between any two points in the network flows along a pre-determined path, connection-oriented technologies make the network more predictable and more manageable. The combination of these attributes helps explain why large networks have been built with a switching fabric in the core of the network (figure 3).

Combining Routing and Switching

While switching technologies have taken over the core of the network, IP routing continues to dominate the edge of the network. The need to

interconnect these two disparate technologies has given rise to the use of overlay networks where the access technology (IP) has been overlaid on top of the core technology (ATM or frame relay). This concept is shown in figure 4. The overlay model extends some of the benefits of the switching technologies to the overall network. Since paths between routers now go through switches, they require a connection, thus bringing more predictability and manageability to the network.

Figure 4

The creation of the overlay model (figure 4) dramatically impacted the operation of the routed portion of the network. When an IP network is overlaid on top of a switched network like ATM, all of the routers appear to be directly connected to each other at the network layer. Therefore, the overlay model requires each router to have an adjacency with every other router in the network. Since the adjacencies must now be established via connections (e.g., ATM VCs), the network now requires a full mesh of VCs to interconnect the routers (figure 5). As the number of routers grows, the number of required VCs grows at the rate of n(n-1)/2, generally referred to as the n^2 problem. The result is a network with a large number of VCs that has scalability problems and becomes very difficult to manage.

The example in figure 5 shows that an overlay network comprised of four routers requires six VCs. However, when a fifth router is added, the VC requirement jumps to ten. As the network grows and the number of routers increases, the VC requirement grows exponentially, thus limiting the scalability of the overall network.
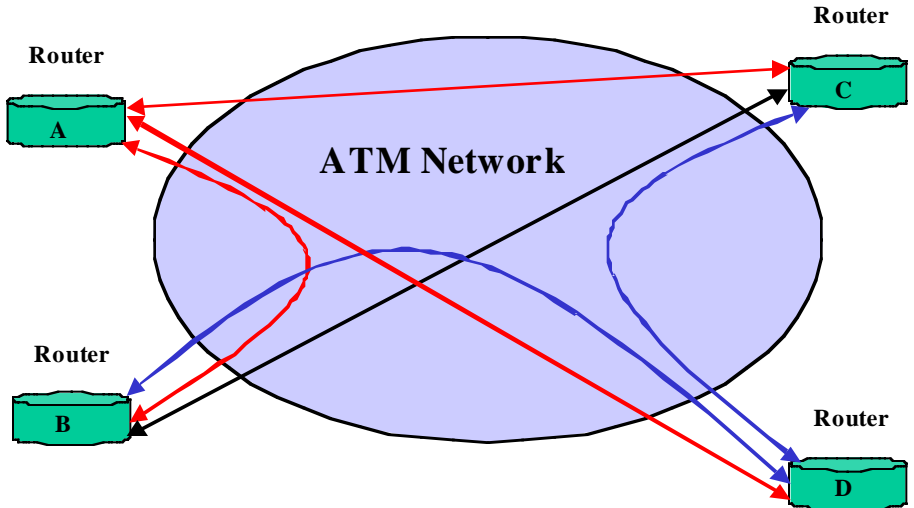
Figure 5

Challenges For Large-Scale Networks

The evolution of telecommunications over the last 100 years has led to an infrastructure embodying the concept of 99.999% uptime, which equates to about 5 minutes of downtime per year. Equipment meeting such stringent requirements is often referred to as being "carrier-class". These attributes are expressed in a number of ways, which are key to service provider objectives, network economics, and rich service offerings.

Scalability

With the rapid growth in bandwidth requirements driven by the increasing use of application rich PCs, modern networks need to scale to an almost infinite capacity. The current technique for running IP over ATM/frame relay networks, the overlay model (figure 4), has proven to have limitations from both a technical perspective, the impact on IP routing protocols, and from an administrative viewpoint, the overhead of managing a large number of virtual circuits.

Robustness

The most significant aspect of current telecommunications networks, particularly the Public Switched Telephone Network (PSTN), is their reliability. Enterprises, and the public at large, have come to expect that they can always pick up the receiver and hear a dialtone. The public data networks of today need to evolve considerably before they can meet the same level of performance.

Feature Richness

Service providers have traditionally delivered rich feature sets by installing specialized networks to meet the specific requirements of a service. With the competitive drive to reduce costs, increase overall efficiency, and introduce new services in combination with existing offerings, unique challenges are posed. The consensus is that packet-based networks, specifically IP, on an end-to-end basis, are best suited to meet these requirements. However, IP networks themselves have traditionally serviced data traffic where a best effort delivery mechanism is sufficient. Transporting deterministic traffic, such as voice, over IP networks requires an evolution in the mechanisms used to route and transport IP.

Network Evolution

The investment made in current networks by service providers and enterprises, both in terms of capital equipment and operational tuning dictates that new services and network growth must be accomplished with minimal network disruption. This includes introducing deterministic services into a non-deterministic IP network, allowing multiple IP traffic types to be established, and enabling the creation and management of IP-based Virtual Private Networks (VPNs).

MPLS Solution

MPLS is designed to meet all the mandatory characteristics of large-scale carrier-class networks. It is evolutionary in the sense that it uses existing layer 3 routing protocols as well as all the widely available layer 2 transport mechanisms and protocols, such as ATM, frame relay, leased lines/PPP and Ethernet. For large-scale public networks, frame relay and particularly ATM are of great interest, primarily because they support underlying concepts of QoS/CoS.

MPLS solves the problem of how to integrate the best attributes of traditional layer 2 and layer 3 technologies. It does this by defining a new operating methodology for the network. The key component within an MPLS network, the label switching router, is capable of understanding and participating in both IP routing and layer 2 switching. By combining these technologies into a singular operating methodology, MPLS avoids the problems associated with methods that define a way for layer 2 and layer 3 to interoperate while maintaining two distinct operating paradigms.

Even though MPLS requires label switching routers (LSRs) to participate in IP routing, the forwarding aspect of MPLS differs significantly from hop-by-hop routing. The LSRs participate in IP routing to understand the topology of the network from a layer 3 perspective. The routing knowledge is then utilized to assign labels to packets. Labels are analogous to the VPI/VCI used in ATM and the DLCI used in frame relay. When viewed on an end-to-end basis, the labels combine to define paths between endpoints. These paths, called label switched paths (LSPs), are intentionally very similar to the connections (VCs) utilized by switching technologies because they provide benefits such as predictability and manageability. In addition, the connection or label switched path (LSP) enables a layer 2 forwarding mechanism to be utilized. As described earlier, a label-swapping mechanism is typically very fast and can be implemented very cheaply in hardware.

Summary

MPLS provides benefits that service providers desperately need in their networks, such as predictability, scalability, and manageability. In addition, work is currently going on that will extend MPLS to provide traffic engineering capabilities for service providers' networks that will enable the service providers to offer differentiated services. Although MPLS will require modifications to existing equipment, it will not require extensive equipment replacement (i.e., no forklift upgrades). Overall, MPLS defines an evolutionary networking paradigm that combines the operating principles of layer 2 and layer 3 technologies while preserving service providers' investment in IP routing technology at the edge of the network and switching technology in the core of the network.

MPLS, like many other technologies, will evolve from its current draft status to a stable standard. However, the road will most likely be very bumpy. In order to be successful, equipment manufacturers that supply the service provider market need to focus on the strategic advantage of their equipment. Even though MPLS will be instrumental in the success of their equipment, manufacturers typically don't have the resources that are required to learn a new technology and keep abreast of that technology as it evolves. That's why many manufacturers are turning to Harris & Jeffries Inc. (H&J) for a solution that will give them the advantage they need to meet their time-to-market goals.

About Harris & Jeffries

Harris & Jeffries, Inc. is the industry's 'ultimate' source for networking software solutions and services. Its *Soft-ATM™*, *High-Performance Frame Relay,™* and *UltraLinq™* product lines provide comprehensive "carrier-class" Asynchronous Transfer Mode, frame relay, redundant, and interworking software capabilities, respectively. H&J's source code enables networking manufacturers, OEMs, and integrators to achieve fast time-to-market with technically advanced products that provide market-leading performance. H&J's products reduce development costs and complexity and are in use by over 120 major networking companies worldwide. H&J can be reached at (781) 329-3200, or at www.hjinc.com.