POINT OF VIEW

MPLS – A Strategic Technology

Executive Summary

Multiprotocol Label Switching (MPLS) is a comprehensive data networking technology that provides many benefits to enterprises and carriers. AT&T has had many years of experience with MPLS, and was an early adopter, announcing its first MPLS-based service in 1999. Since then, AT&T has continually rolled out new and enhanced MPLS-based IP VPN services in support of enterprise customers. Today, AT&T is regarded by leading telecommunications analysts as having one of the most comprehensive VPN portfolios in the industry, including MPLS, IPSec and SSL-based solutions.

Now considered a mainstream technology throughout the telecommunications industry, MPLS is the key technological component underpinning AT&T's current, and future, network evolution. AT&T has adopted MPLS as a strategic platform onto which all of its diverse data networks will converge into a single, seamless global MPLS network, transported on an intelligent optical infrastructure.

POINT OF VIEW

Introduction

Multiprotocol Label Switching (MPLS) is a standardized protocol and comprehensive unifying networking architecture. The design of MPLS follows the principles described by the classic 1984 paper on systems design, "End-to-End Arguments in System Design," by J.H. Saltzer, D.P. Reed and D.D. Clark. In this seminal work, the authors describe the principles of using simplicity in the core of a system and complexity on the edges. To both simplify and increase the efficiency of core transport, the MPLS protocol enables data to be transmitted efficiently across a network infrastructure utilizing a technology known as "label switching." As customers' data enters the network, a 20 bit header called a label is appended to each packet. Labels can be used to convey several types of information about a packet, but probably the most frequent use of a label is to uniquely identify a customer's Virtual Private Network in a shared infrastructure and keep it private. Used in this fashion, a label uniquely identifies a packet as belonging to a specific IP VPN. Upon reaching its destination, the label is removed, thereby returning the data packet to its original state. The process is seamless and unnoticeable to end-users. One can think of MPLS in this context as a "special delivery courier service" for network data packets.

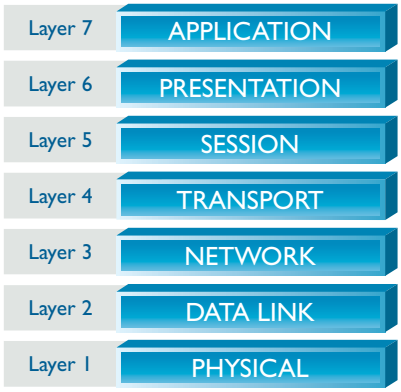

The Seven Layer OSI Model

There has been much confusion in the industry regarding terminology related to VPNs and MPLS. In this paper we'll be using the following definitions. "MPLS VPN" will be used to describe the full range of VPN functionality supported by MPLS: both Layer 2 and Layer 3 VPN services over an MPLS core. This comprehensive and somewhat generic term includes Frame Relay, Asynchronous Transfer Mode (ATM) *and* IP VPNs. An "IP VPN" is a Layer 3 IP-routed service which is provided to the customer; various technologies can be used to create an IP VPN, such as IPSec, MPLS, SSL, etc. This paper does not describe all possible IP VPN technologies but is focused on MPLS, the technology, its key attributes, and how it can be used to create IP VPNs. The term "IP VPN" does not describe the type of network upon which the service is provisioned; it is possible to create IP VPNs on both public (connected to the Internet) and private (not connected to the Internet) IP networks. Lastly, an "MPLS-based IP VPN" is an IP VPN which is provisioned over a network which is MPLS-enabled. Within the context of any specific paragraph where one of these terms is used, "VPN" without any modifiers may be used later in the same paragraph after the specific type of VPN has been identified.

Industry Perspective on MPLS

Technology research firm, Gartner, Inc., states, "MPLS is more than hype; this is the next-generation enterprise WAN."

Analysts, vendors and carriers alike are touting MPLS' benefits and predicting wide and almost universal deployment. Technology research firm, Gartner, Inc., states, "MPLS is more than hype; this is the next-generation enterprise WAN."[1] They predict that by 2006 MPLS "will have evolved solidly into the industry's preferred technology/ service for replacing frame relay, private line and, to a lesser degree, ATM service."[2] In late 2002, Cisco Systems stated that over 100 carriers were known to have deployed MPLS in their networks, clearly indicating that it is overwhelmingly considered a proven and stable technology.

Perhaps the most attention-grabbing application of MPLS has been its use as a key enabler of network-based IP Virtual Private Networks. In-Stat/MDR says, "IP VPNs are probably the hottest area in wide area networking today."[3] Enterprises are turning to IP VPNs for a wide range of reasons, including reduced total cost of ownership, flexible architecture, support of extranet environment, meshed configurations and future-proof architecture. Within the last two years, MPLS has risen to prominence as the primary technological facilitator of IP VPNs. Most analyst reports on IP VPNs now discuss MPLS as a matter of course. Their view is that an MPLS-based IP VPN is mandatory for carriers who wish to be regarded as having a full suite of IP VPN offerings.

Carriers have found that the use of MPLS provides additional opportunities for service innovation and enhancements. MPLS has the ability to provide Quality of Service/ Class of Service (QoS/CoS) and to facilitate network Service Level Agreement guarantees. It should not be a surprise that the earliest market entries for providing voice and multimedia services over an IP infrastructure have been implemented using MPLS-based VPNs with the requisite QoS.

MPLS-What it is and How it Works

MPLS is one of the most innovative data networking technologies to emerge since the rise of the Internet. A simple indexing mechanism called a "label" replaces traditional IP packet forwarding, where complicated address matching is performed at each hop in the network. The label describes how the packet should be handled within the network and thus assigns the packet to a Forwarding Equivalence Class (FEC). As a packet traverses the network the intermediate nodes simply swap labels and forwards the packet based on the FEC, without ever examining the contents of the packet. Thus all packets which belong to the same FEC get treated in the same way and quickly are sped along their way. Label-swapping is considered to be more like ATM switching in its speed and simplicity.

Packets are forwarded along a "label switched path (LSP)", where each "label switch router (LSR)" makes forwarding decisions based solely on the contents of the label. At each hop, the LSR strips off the existing label and applies a new label that tells the next hop LSR how to forward the packet. Labels are distributed between LERs and LSRs using the "label distribution protocol" (LDP). Label Switch Routers in an MPLS network regularly exchange label and reachability information with each other using standardized procedures in order to build a complete picture of the network they can then use to forward packets.

Label Switch Paths (LSPs) are established by the network operator for a variety of purposes, such as to create network-based IP Virtual Private Networks or to route traffic along specified paths through the network. In many respects, LSPs are no different than PVCs in ATM or Frame Relay networks, except that they are not dependent on a particular Layer 2 technology.

Consistently an industry leader in the use of MPLS, AT&T's most recent and significant announcement regarding MPLS to date took place in early 2003. In a briefing to key industry analysts and media, AT&T announced the use of MPLS in support of its overall long-term network convergence plans.

In response, Broadband Publishing Corporation, in its Network Technology Report, described this announcement as "the strongest endorsement of MPLS to date, from the largest and most financially stable network operator in North America."[4] Broadband Publishing went on to predict that AT&T's endorsement of MPLS may be a harbinger of the future technology direction of the industry, as it was ten years ago when AT&T announced its support of frame relay, now a widely deployed mainstream networking technology.

AT&T believes that adoption of Multiprotocol Label Switching (MPLS) technology is *the* overriding telecommunications strategic direction of the industry. Speaking at the announcement of the merger of the Frame

POINT OF VIEW

Relay and MPLS Forums, recently retired AT&T Network Services President Frank Ianna described AT&T's perspective on MPLS as a strategic technology for carriers: "We envision a future network that will replace today's multiple networks (ATM, Frame Relay, Private Line) with a single, global, MPLS-enabled backbone over an intelligent optical IP-based core, with intelligent nodes and multi-protocol/multi-service capabilities at its edges."

AT&T believes that adoption of Multiprotocol Label Switching (MPLS) technology is the overriding telecommunications strategic direction of the industry.

AT&T's MPLS Key Milestones

- **1Q98** Announced MPLS VPN Service: AT&T IP Enabled Frame Relay (IP FR)
- **2Q99** First Production Customer IP FR
- **3Q01** Announced COS: IP FR
- **2Q02** Announced End to End SLAs-IP FR
- **2Q02** Announced Availability-MPLS IP VPN for Service Providers
- **1Q03** Announced GA: IP FR on AT&T Global Network
- **1Q03** Announced GA fully managed IP VPN including COS: AT&T Enhanced VPN

Key Benefits of MPLS

MPLS could be said to combine the best features of Layer 2 and Layer 3 networking technologies in a single architecture, thus bringing significant benefits to both enterprises and carriers. To traditional IP networks, MPLS adds the significant reliability and performance capabilities for which Layer 2 networks are known, in addition to a hierarchical network design which allows scaling beyond the capabilities of a meshed Layer 2 design. Most importantly, MPLS is a key enabler of important IP based services, such as IP VPNs and Class of Service. Because MPLS provides *multiprotocol* transport, it can also be used to transport Layer 2 services, such as Frame Relay, ATM and Ethernet. In a large enterprise, this provides the ability for virtually any combination of endpoints to be interconnected via any technology within the same MPLS VPN, simply and flexibly. Moreover, the interconnection is secure, and service can be delivered with performance guarantees. For carriers, this potentially enables network convergence as both Layer 2 and 3 services can be reliably transported across the same network.

Reliability is one of the most important attributes of any network architecture, both to customers and carriers. As businesses become more heavily dependent upon networking as an integral part of their core business, high reliability is imperative. MPLS adds significant reliability and rerouting capabilities to an IP network through its Fast Reroute (FRR) feature. This feature enables rapid rerouting of traffic around a failed network component such as an interface, line card or link.

One typical use of FRR is for a carrier to establish a backup path for each and every link or path in the network, similar to the physical layer restoration provided by SONET rings, only with many more restoration options and flexibility. In case of a physical link failure, the backup path can be automatically enabled to immediately allow the traffic to resume via the designated "restoration" path. The restoration time is sub-second – significantly faster than the usual time it takes for routing protocols to re-converge after a link failure in today's traditional IP network. With this sort of reliability and instant fallback, carriers can provide not only best-effort but also premium services to support latency-sensitive applications such as voice or video across an MPLS network.

Class of Service/Quality of Service (CoS/QoS) for Application Support

While high reliability is crucial to enterprises' networking environments, they are also increasingly looking for Class of Service/Quality of Service support to provide differentiated performance for distinct applications across a converged network. MPLS' basic design and traffic engineering capabilities inherently provide the framework to support Class of Service/Quality of Service (CoS/QoS). How does MPLS do this? At the edge of the network Differentiated Services (DiffServ[5]) packet classifications are mapped to MPLS labels through the use of the Experimental field in the MPLS header. This field is used to specify the per-hop behavior of the packet within the network, such as scheduling and drop preference parameters. DiffServ-marked packets are distinguishable when carried across an MPLS backbone and thus receive the correct performance-handling by the network

Using this capability, an enterprises' best-effort traffic, such as email, and delay-sensitive applications, such as voice and video, can be carried across the same network as ordinary data traffic, with each type of application being given the appropriate handling and priority via MPLS. This gives customers a cost-effective approach to manage bandwidth in order to achieve specific performance levels.

Layer 2 Support for Network Evolution

Another key MPLS benefit is its encapsulated support of Layer 2 protocols across a Label Switched Path. This attribute has two aspects – the ability to run Layer 2 services "under" MPLS and the ability to enable an easy transition from a traditional Layer 2 service to an IP VPN.

Ethernet, Frame Relay and ATM can run "under" any service, as a point-to-point Layer 2 VPN service (according to IETF Martini or Kompella draft implementations) or as a multipoint Layer 2 VPN service (transparent LAN, or "Ethernet over MPLS" using IETF Virtual Private LAN Service drafts). These Layer 2 MPLS VPN 'flavors' create secure but unencrypted switching paths across the service provider's private network and have emerged to challenge IPSec VPNs because of their ease of deployment and management from a customer's perspective – no special equipment to manage, simply connect to the carrier's MPLS network.

Support for Layer 2 protocols also enables customers to smoothly transition from a traditional Layer 2 network, such as frame relay and ATM, to an MPLS-based IP VPN without making expensive and disruptive changes in their equipment or addressing. For carriers, Layer 2 support provides a platform for network convergence as it enables legacy Layer 2 services to be transported across an MPLS cloud.

IP VPN Support

The most significant direct customer benefit of MPLS is as an enabling platform for creating secure, network-based IP VPNs.

The most significant direct customer benefit of MPLS is as an enabling platform for creating secure, network-based IP VPNs. MPLS-based IP VPNs, as described in IETF standard RFC2547bis[6], combine the security, performance and reliability of legacy data protocols such as Frame Relay and ATM with the routing flexibility of IP. MPLS-based IP VPNs are access method agnostic, thus Frame Relay, ATM, DSL, dial, wireless, cable and Ethernet can all be used for access. This flexibility simplifies the

creation of a VPN; it is not necessary to redeploy endpoints with a consistent technology in order to create a single VPN, simply connect any possible variety of disparate endpoints to the carrier's MPLS network and the VPN is in service. It also makes it possible for network administrators to respond to mergers and acquisitions without requiring reinstallation of network endpoints; simply add a connection to the corporate VPN.

Any to Any Connectivity for Network Scalability. As enterprises become a collection of different endpoints with disparate application needs, the model of connecting to the headquarters mainframe becomes less the norm; offices need access to distributed applications located on servers throughout an organization. Without a full mesh of interconnections, delay is created as traffic tandem routes through a hub location.

An MPLS-based IP VPN solves this problem because it creates a private, meshed network of the customer's locations. An MPLS-based IP VPN solves the proverbial "N squared"[7] scaling problem of Layer 2 networks such as Frame Relay and ATM where many dedicated permanent virtual circuits (PVCs)/Virtual Circuits(VCs) are required if all customer locations must communicate with each other.

With an MPLS-based IP VPN, authorized customer locations are provided with VPN connectivity to *all* other authorized customer sites within the same VPN in a completely meshed fashion, similar to a private Internet in operation. But unlike the Internet, an MPLS-based IP VPN is as secure as frame relay or ATM, and is designed to prevent traffic from unauthorized sources from accessing a customer's VPN.

Simplified Network Integration/Migration through Flexible Addressing Support. An MPLS-based IP VPN supports any type of customer IP addressing, thereby providing maximum flexibility to enterprises – renumbering is not required when connecting to a carrier's MPLS network. Precisely because of the privacy of the VPN; the customer's IP addresses are only visible within that VPN and not beyond. Customers are allowed to use any IP addressing scheme they wish, as long as uniqueness is maintained within the customer's entire VPN. For a large enterprise, this potentially means there is no need to renumber internal networks at far-flung offices when establishing a VPN between them. This numbering flexibility can also ease the technical difficulties associated with merging data networks when mergers and acquisitions take place since massive network renumbering is not required.

Connecting to an MPLS-based IP VPN service is simple for an enterprise customer.

Routing Simplicity. Connecting to an MPLS-based IP VPN service is simple for an enterprise customer. Customers need not run MPLS within their networks in order to take advantage of an MPLS-based IP VPN; MPLS is managed completely by the carrier, giving enterprises the benefits of a secure, scalable, virtual private network without administrative complexity. Either static routing or BGP4[8] routing protocols are supported between the carrier's network equipment (called the Provider Edge or PE, in MPLS terminology) and the customer's router (called the Customer Edge or CE). No specialized routing configurations are required and many carriers, including AT&T, will also manage the CPE. This saves time and effort on the part of the customer's IT or engineering staff.

Class of Service Support. As described previously, MPLS supports Class of Service, thus deployment of an MPLS-based IP VPN enables all of an enterprise's applications to be carried across the same network. This gives customers the capability to obtain different performance characteristics for key applications without requiring the use and expense of multiple networks.

Traffic Engineering and Scalable, Hierarchical Architecture

> ***MPLS brings multiple benefits to carriers, and thus indirectly to customers, through its traffic engineering capabilities and scalable, hierarchical design.***

MPLS brings multiple benefits to carriers, and thus indirectly to customers, through its traffic engineering capabilities and scalable, hierarchical design. MPLS traffic engineering, accomplished via explicit routing, gives carriers an important tool for managing their networks to provide peak performance as a significant augment to traditional capacity planning.

In traditional IP networks without MPLS, typically IP traffic is routed via a designated Interior Gateway Protocol (IGP), such as OSPF, IS-IS, RIP, among others. Routing is done on a hop by hop basis, determined by a set of pre-determined criteria; it's not possible to decide in advance how to route a given traffic flow. Each individual hop in the network calculates the shortest path to the next hop, and the result of the calculation is always the same. Therefore, even if there is congestion on a particular link, packets will still be routed via the congested link if the routing protocol thinks that link is the "shortest path."

In an MPLS network, a Label Switched Path (LSP) functions like a virtual circuit in a Layer 2 network, adding connection-oriented functionality onto a connectionless IP network. MPLS adds a rich set of routing metrics that are more extensive than those of ordinary IP routing protocols. These metrics enable the switches to make routing decisions moment by moment based on a robust set of parameters such as bandwidth, class of service and delay, rather than just "shortest path." Unlike an ordinary IP network, MPLS traffic engineering allows the carrier to specify a particular route for an LSP which is different than the path selected by the interior routing protocol (such as OSPF) using its metrics. By creating an LSP for traffic engineering purposes, the carrier can thus route traffic via a predetermined, specified path based on the carrier's criteria, such as to create paths providing stringent performance requirements to meet a customer's need. A carrier may also use traffic engineering as an additional tool in its capacity management arsenal, such as by creating an LSP to avoid a frequently congested route until more capacity can be deployed. In another example, a carrier may choose to deliberately carry large volumes of traffic across a very high-speed link to obtain the most efficient use of deployed network capacity. Traffic engineering can also be used as a tool to help deliver end-to-end Class of Service and thus to give certain types of traffic priority in their handling and routing throughout the network.

MPLS technology brings one additional, important benefit to carriers, from which customers derive considerable advantage: the ability to create a hierarchical and scalable network architecture that grows and evolves with changing customer requirements. By separating the functionality and complexity required at the core from that required at the edges of the network, scalability issues that might impede network growth are minimized. A scalable network enables a carrier to develop and provide additional services over time as customer requirements

evolve, without requiring disruption of the customers' installed networks. Scalability also means that customers' networks are not disturbed and put out of service as the carrier's network grows, thus ensuring continuous, reliable operation even as the network evolves. In an MPLS network, the edges maintain the customer and Internet intelligence while the core provides for fast, reliable, secure and congestion-free transport. It is this hierarchical structure, made possible by MPLS technology, that allows a carrier to securely provide Internet services and VPN services over the same core network.

AT&T's Perspective on MPLS

AT&T has been deeply involved with MPLS since its inception and is regarded as one of the industry leaders based on its early and continuing work with this technology.

AT&T has been deeply involved with MPLS since its inception and is regarded as one of the industry leaders based on its early and continuing work with this technology. One visible example of AT&T's involvement with MPLS is that AT&T Labs engineers are listed as co-authors of the IETF standard RFC2547bis. This IETF document describes a method for creating private IP VPNs across a shared network infrastructure through the use of MPLS and is considered an industry standard. AT&T was the first carrier in the United States to offer a service based on this then-nascent technology, announcing its Enterprise Class Virtual Private Network (VPN) Services on January 26, 1999. Since 1999, AT&T has steadily rolled out new MPLS-based services or enhancements, the most recent being its global, fully-managed Enhanced VPN Service, announced in early 2003.

Despite the complexity of upgrading large networks carrying massive volumes of traffic, and with many thousands of customer connections, AT&T has already added MPLS functionality to all of its data networks, all the while providing highly reliable service that is AT&T's trademark.

AT&T Labs MPLS Industry Technical Leadership

AT&T Labs engineers are involved in many facets of MPLS' design and evolution in the Internet Engineering Task Force (IETF) and International Telecommunications Union (ITU). AT&T Labs engineers are actively working as co-authors in the following areas of development for MPLS:

- **RFC2547bis (BGP/MPLS VPN)**
- **MPLS NM/OAM/MIBs**
- **MPLS Traffic Data**
- **MPLS Protection/Restoration**
- **MPLS/DiffServ Traffic Engineering**
- **VoIP over MPLS Header Compression**
- **OSPF Congestion Control/Failure Recovery**

MPLS and AT&T's Strategic Network and Services Evolution

MPLS is a fundamental architectural element in AT&T's network and services evolution. AT&T is evolving all of its networks to a single, seamless next generation network, enabled by MPLS, transported across an intelligent optical infrastructure. It will be an entirely new self-operating network providing customers with high performance, high reliability and flexibility. The core of the network will be MPLS over an intelligent optical core, and will be managed to provide extremely high reliability in keeping with AT&T's very low Defects per Million[9] network reliability measurement standards.

In addition to building a new, high-performance core network, AT&T is also investing at the "edges," to create a sophisticated e-bonding environment where systems talk to systems, eliminating the possibility of human error. The end result will be the creation of the industry's best customer experience, from ordering to billing, with extremely high quality and reliability.

AT&T's network evolution will create a highly reliable, high-performance, secure network that will support all services and technologies, enabling the interconnection of any possible type of customer endpoints in an integrated fashion. This will enable enterprises to support their business applications and achieve the flexible, integrated networking which will help them achieve higher productivity and efficiency.

Summary

MPLS is a rich, full-featured networking technology that provides benefits to enterprises and carriers alike.

MPLS is a rich, full-featured networking technology that provides benefits to enterprises and carriers alike. Enterprises benefit from MPLS directly without having to run it in their own networks; MPLS enables the creation of secure, reliable VPNs which are simple to manage, easy to deploy and which provide Class of Service/Quality of Service support. AT&T has taken a leadership position in exploiting MPLS' capabilities to create MPLS-based services and is regarded as a market leader for its IP VPN portfolio. AT&T is using MPLS as a strategic convergence platform that enables the integration of its disparate networks into a seamless global MPLS network, supported by an intelligent optical infrastructure. AT&T will continue to create and offer innovative services that meet the enterprise networking needs of its business customers, now and in the future.

Sources

1. Pultz, J. and Rickard, N., "MPLS Networks: Drivers Beat Inhibitors in 2003," Gartner Research, 10 February 2003

2. Carr, Charles, "IP VPN: Hitting the Big Time," Gartner Research, 20 January 2003

3. Goldberg, Henry, "End-User Demand and Perspectives on IP VPNs", In-Stat/MDR, January 2003

4. Broadband Publishing, "Network Technology Report," ISSN 1542-6009, 2003, page 1.

5. DiffServ is an abbreviation for Differentiated Services. The Differentiated Services architecture is a set of IETF standards which describe an architecture for achieving class of service functionality within an IP network. An excellent explanation of DiffServ is found in Cisco's white paper located at: http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/difse_wp.pdf

6. http://www.ietf.org/internet-drafts/draft-ietf-ppvpn-rfc2547bis-03.txt

7. The "N squared problem" in networking is based on a fundamental law of graph theory which states that the number of interconnections increases by the square of the number of nodes on the network. Thus, connection-oriented network such as in a private line, frame relay or ATM where a fully meshed network is required, adding one new node requires adding connections to all other existing nodes to maintain a full mesh.

8. BGP4 stands for Border Gateway Protocol, Version 4. BGP4 is an exterior routing protocol which routes traffic between autonomous systems. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).

9. AT&T measures the reliability of its various networks using the rigorous Defects Per Million (DPM) methodology, where outages are measured from the perspective of a customer experiencing an outage. Thus, for IP, 1 minute of a customer port outage counts as a defect, as does 1 minute of PVC outage for frame relay.