CHAPTER 1

Introduction to MPLS

This chapter is an overview of Multiprotocol Label Switching (MPLS), highlighting MPLS in ATM networks and packet-based networks. It concentrates on the fundamentals of MPLS network design that apply to all ATM MPLS networks, including those supporting VPNs and traffic engineering.

- What is MPLS?
- Label Switching Features
- Label Switching Benefits
- IMPLS Compared to Other IP-over-ATM Schemes
- MPLS Network Structure
- MPLS Applications
- MPLS Virtual Private Network
- References

What is MPLS?

Multiprotocol Label Switching (MPLS) is a high-performance method for forwarding packets (frames) through a network. It enables routers at the edge of a network to apply simple labels to packets (frames). ATM switches or existing routers in the network core can switch packets according to the labels with minimal lookup overhead.

The BPX® 8650 is an IP+ATM switch that provides ATM-based broadband services and integrates Cisco IOS® software via Cisco 7200 series routers to deliver Multiprotocol Label Switching (MPLS) services.

MPLS integrates the performance and traffic management capabilities of Data Link Layer 2 with the scalability and flexibility of Network Layer 3 routing. It is applicable to networks using any Layer 2 switching, but has particular advantages when applied to ATM networks. It integrates IP routing with ATM switching to offer scalable IP-over-ATM networks.

In contrast to label switching, conventional Layer 3 IP routing is based on the exchange of network reachability information. As a packet traverses the network, each router extracts all the information relevant to forwarding from the Layer 3 header. This information is then used as an index for a routing table lookup to determine the packet's next hop. This is repeated at each router across a network. At each hop in the network, the optimal forwarding of a packet must be again determined.

The information in IP packets, such as information on IP Precedence and Virtual Private Network membership, is usually not considered when forwarding packets. Thus, to get maximum forwarding performance, typically only the destination address is considered. However, because other fields could be relevant, a complex header analysis must be done at each router that the packet meets.

The main concept of MPLS is to include a *label* on each packet.

Packets or cells are assigned short, fixed length labels. Switching entities perform table lookups based on these simple labels to determine where data should be forwarded.

The label summarizes essential information about routing the packet:

- Destination
- Precedence
- Virtual Private Network membership
- Quality of Service (QoS) information from RSVP
- The route for the packet, as chosen by traffic engineering (TE)

With Label Switching the complete analysis of the Layer 3 header is performed only once: at the edge label switch router (LSR), which is located at each edge of the network. At this location, the Layer 3 header is mapped into a fixed length label, called a label.

At each router across the network, only the label need be examined in the incoming cell or packet in order to send the cell or packet on its way across the network. At the other end of the network, an edge LSR swaps the label out for the appropriate header data linked to that label.

A key result of this arrangement is that forwarding decisions based on some or all of these different sources of information can be achieved by means of a single table lookup from a fixed-length label. For this reason, label switching makes it feasible for routers and switches to make forwarding decisions based upon multiple destination addresses.

Label switching integrates switching and routing functions, combining the reachability information provided by the router function, plus the traffic engineering benefits achieved by the optimizing capabilities of switches. These benefits are described in more detail in the next section.

Label Switching Features

MPLS, in conjunction with other standard technologies, offers many features critical for service providers:

- MPLS, in combination with the standard IP routing protocols OSPF or IS-IS, provides full, highly scalable support of IP routing within an ATM infrastructure.
- MPLS, in combination with the Border Gateway Protocol (BGP), provides support for highly scalable IP Virtual Private Network (VPN) services. IP VPN services are an invaluable development in provider networks, giving enterprise customers a service that meets their needs for private, connectionless delivery of IP services.
- Service-Level Agreements may be provided in a form suitable for connectionless traffic. Cisco networks assist the process of providing Service-Level Agreements by supporting MPLS in combination with forthcoming DiffServ standard. Along with supporting Virtual Private Networks, the ability to offer Service Level Agreements suitable for IP traffic is a critical requirement to meet new demand for IP services.
- Cisco’s implementation of MPLS allows support for harder quality-of-service where required, using full ATM switch capabilities.

Cisco IP+ATM networks fully support all relevant IP routing protocols and MPLS, while fully supporting traditional ATM services. MPLS and IP routing can readily be introduced to traditional ATM networks by using PVP or PVC tunnels, as MPLS-capable switches are continuously introduced.

Cisco IP+ATM switches allow carriers to continue to meet their existing demand for virtual circuit services while adding optimized support for critically important new services: IP and IP Virtual Private Networks. Furthermore, Cisco supports all of the standards relevant to carrier-class IP services: MPLS, the Multiprotocol Border Gateway Protocol, other standard routing protocols, and MPLS Traffic Engineering.

Label Switching Benefits

MPLS offers many advantages over traditional IP over ATM.

When integrated with ATM switches, label switching takes advantage of switch hardware optimized to take advantage of the fixed length of ATM cells and to switch the cells at high speeds. For multiservice networks, label switching enables the BPX switch to provide ATM, Frame Relay, and IP Internet service all on a single platform in a highly scalable way. Support of all these services on a common platform provides operational cost savings and simplifies provisioning for multiservice providers.

For internet service providers (ISPs) using ATM switches at the core of their networks, label switching enables the Cisco BPX 8600 series, the 8540 Multiservice Switch Router, and other Cisco ATM switches to provide a more scalable and manageable networking solution than overlaying IP over an ATM network. Label switching avoids the scalability problem of too many router peers and provides support for a hierarchical structure within an ISPs network.

These MPLS benefits are analyzed in greater detail:

- **Integration**
  When applied to ATM, MPLS integrates IP and ATM functionality rather than overlaying IP on ATM. This makes the ATM infrastructure visible to IP routing and removes the need for approximate mappings between IP and ATM features. MPLS does not need ATM addressing and routing techniques such as PNNI, although these can be used in parallel if required.
- **Higher Reliability**
  In Wide Area Networks (WANs) with ATM infrastructures, MPLS is an easy solution for integrating routed protocols with ATM. Traditional IP over ATM involves setting up a mesh of Permanent Virtual Circuits (PVCs) between routers around an ATM cloud, and the Next Hop Resolution Protocol (NHRP) achieves a similar result with switched virtual circuits (SVCs). But there are a number of problems with this approach, all arising from the method that the PVC links between routers are overlaid on the ATM network. This makes the ATM network structure invisible to the routers. A single ATM link failure could make several router-to-router links fail, creating problems with large amounts of routing update traffic and subsequent processing. (See The Problem of Persistent Loops Due to Protocol Conflicts, page 1-5)
- **Better Efficiency**
  Without extensive tuning of routing weights, all PVCs are seen by IP routing as single-hop paths with the same cost. This might lead to inefficient routing in the ATM network.
- **Direct Classes of Service Implementation**
  When used with ATM hardware, MPLS makes use of the ATM queueing and buffering capabilities to provide different Classes of Service (CoS). This allows direct support of IP Precedence and CoS on ATM switches without complex translations to the ATM Forum Service Classes.

- **More Elegant Support of Multicast and RSVP**
  In contrast to MPLS, overlaying IP on ATM has other disadvantages, particularly in support of advanced IP services such as IP multicast and RSVP. Support of these services entails much time and work in the standards bodies and implementation; the resulting mapping between IP features and ATM features is often approximate.
- **VPN Scalability and Manageability**
  MPLS can make IP Virtual Private Network services highly scalable and very easy to manage. Virtual Private Network services are an important service for providing enterprises with private IP networks within their infrastructures. When an ISP offers a VPN service, the carrier supports many individual VPNs on a single infrastructure. With an MPLS backbone, VPN information can be processed only at the ingress and exit points, with MPLS labels carrying packets across a shared backbone to their correct exit point. In addition to MPLS, the Multiprotocol Border Gateway Protocol (BGP) is used to deal with information about the VPNs. The combination of MPLS and Multiprotocol BGP makes MPLS-based VPN services easier to manage, with straightforward operations to manage VPN sites and VPN membership. It also makes MPLS-based VPN services extremely scalable, with one network able to support hundreds of thousands of VPNs.
- **Reduces Load on Network Cores; More Robust**
  VPN services demonstrate how MPLS supports a hierarchy of routing knowledge. Additionally, you can isolate Internet routing tables from service provider network cores. Like VPN data, MPLS allows access to the Internet routing table only at the ingress and exit points of a service provider network. With MPLS, transit traffic entering at the edge of the provider's autonomous system can be given labels that are associated with specific exit points. As a result, internal transit routers and switches need only process the connectivity with the provider's edge routers, shielding the core devices from the overwhelming routing volume exchanged in the Internet. This separation of interior routes from full Internet routes also provides better fault isolation and improved stability.
- **Traffic Engineering Capabilities**
  Other benefits of MPLS include traffic engineering (TE) capabilities needed for the efficient use of network resources. Traffic engineering enables you to shift the traffic load from overutilized portions to underutilized portions of the network, according to traffic destination, traffic type, traffic load, time of day, and so on.

IMPLS Compared to Other IP-over-ATM Schemes

In ATM networks, MPLS allows ATM switches to directly support IP services, giving maximum efficiency compared to other approaches. Traditional IP over ATM connects routers over Permanent Virtual Circuits (PVC).

Cisco also supports an alternative IP-over-ATM scheme called Multiprotocol Over ATM (MPOA), which uses the Next Hop Resolution Protocol (NHRP). Unlike MPLS, MPOA overlays IP over ATM rather than fully integrating them. Although they do not share many of the advantages of MPLS in the WAN, MPOA and NHRP are cost-effective technologies for interconnecting nearby emulated LANs (ELANs) at high speeds. MPOA and similar proprietary approaches carry IP traffic over Switched Virtual Circuits (SVC). Traditional IP over ATM, MPOA, and proprietary approaches all have similar disadvantages:

- It is difficult to offer some types of IP services on the networks. For example, IP Class of Service cannot be offered natively by traditional ATM switches, and must be offered by translation to quite different ATM Forum Quality of Service concepts.

- Where IP services are offered, they are difficult to administer. Two levels of routing must be administered: IP routing (via OSPF or EIGRP or similar) and PNNI or similar routing for ATM. MPOA requires additional administration. Service translations, for example IP Class of Service to ATM Quality of Service, also require administration.
- IP services can be quite inefficient over ATM networks. For example, IP Multicast over ATM networks is difficult to achieve on a large scale due to the interaction of multicast routing, multicast group membership processing and ATM VC maintenance.
- There can be scaling limitations and/or dangerous interactions between IP routing (OSPF, and so on) and the ATM network, leading to unstable networks. Traditional IP over ATM can lead to storms of IP routing updates and subsequent network meltdown, if more than 30 OSPF routers are connected in a full mesh over PVCs. MPOA is unsafe when connecting routers to each other, and is intended only to connect hosts to routers or hosts to hosts. (See below.)
- IP services require a substantial implementation and management effort. For example, an MPOA implementation requires PNNI, SVC signalling, ATM ARP, an ATM ARP server, NHRP, and a NHRP server, in addition to AAL5, IP routing (OSPF, and so on) and an IPv4 stack.

MPLS in ATM networks avoid all of these disadvantages.

The Problem of Persistent Loops Due to Protocol Conflicts

If *N* number of routers are running OSPF and are connected in a full mesh over ATM PVCs, a single physical ATM link failure may result in ATM-layer rerouting of a large number of PVCs. If this takes too long, or if the ATM network cannot reroute PVCs at all, a large number of PVCs effectively fails.

The number of PVCs involved may be of the same order magnitude as *N*, and even *N*2 in some cases. In any case, it is likely to be seen by O(*N*) routers, where "O(N)" means "a number proportional to *N*". So, a single ATM link failure will cause each of O(N) routers to send a link state advertisement (LSA) of size (at least) O(*N*) to (*N*-1) neighbors. Thus a single event in the ATM network results in O(*N*3) to O(*N*4) traffic.

When a router receives an LSA, it must immediately recalculate its routing table because it must not forward packets based on old routing information. The processor load caused by a storm of routing updates might cause the routers to drop or not send keep-alive packets, which appears to the neighboring routers as further link failures. These lead to further LSAs being sent, which perpetuates the problem.

The net result is that a full mesh network can go persistently unstable after a single network event.

This critical failure occurs because the routers do not see the state of the ATM links and switches directly. IS-IS has somewhat better performance than OSPF in full mesh conditions because IS-IS has more sophisticated flooding capabilities (these capabilities, specifically the ability to pace flooding and block flooding on some interfaces, are also becoming available on OSPF). However this does not address the underlying problem.

The solution is to enable IP routing to directly see the state of ATM links, which is what is done by ATM MPLS.

MPLS addresses the fundamental problem underlying the instability of the full mesh network: the basic conflict between routing protocols. PNNI routing at the ATM layer can make decisions that conflict with OSPF or similar routing at the IP layer. These conflicting decisions can lead to persistent loops. (See the NHRP Protocol Applicability Statement, RFC2333, for more on this. Further investigation on router-to-router NHRP at the IETF revealed that router-to-router NHRP was not practical.)

The only reliable solution to this problem is to use the same routing protocol at the IP layer and ATM layer. This is exactly what MPLS does in ATM networks.

MPLS Network Structure

A typical structure for Multiprotocol Label Switching networks used by providers (carriers or ISPs) is shown in Figure 1-1.

The basic elements in a label switching network are:

- **Edge Label Switch Routers**
  Edge Label Switch Routers are located at the boundaries of a network, performing value-added network layer services and applying labels to packets. These devices can be either routers, such as the Cisco 7500, or multilayer LAN switches, such as the Cisco Catalyst 5000.
- **Label Switches**
  These devices switch labeled packets or cells based on the labels. Label switches may also support full Layer 3 routing or Layer 2 switching in addition to label switching. Examples of label switches include the Cisco 6400, the Cisco 8540 Multiservice Switch Router, Cisco BPX 8650, and Cisco 7500.
- **Label Distribution Protocol**
  The Label Distribution Protocol (LDP) is used in conjunction with standard network layer routing protocols to distribute label information between devices in a label switched network.

An MPLS network consists of Edge Label Switch Routers (Edge LSRs around a core of Label Switch Routers (LSRs). Customer sites are connected to the provider MPLS network.

Typically there are several hundred customer sites per edge LSR. The Customer Premises Equipment (CPE) runs ordinary IP forwarding but usually does not run MPLS. If the CPE does run MPLS, it uses it independently of the provider.

Figure 1-1 Typical MPLS Network Structure

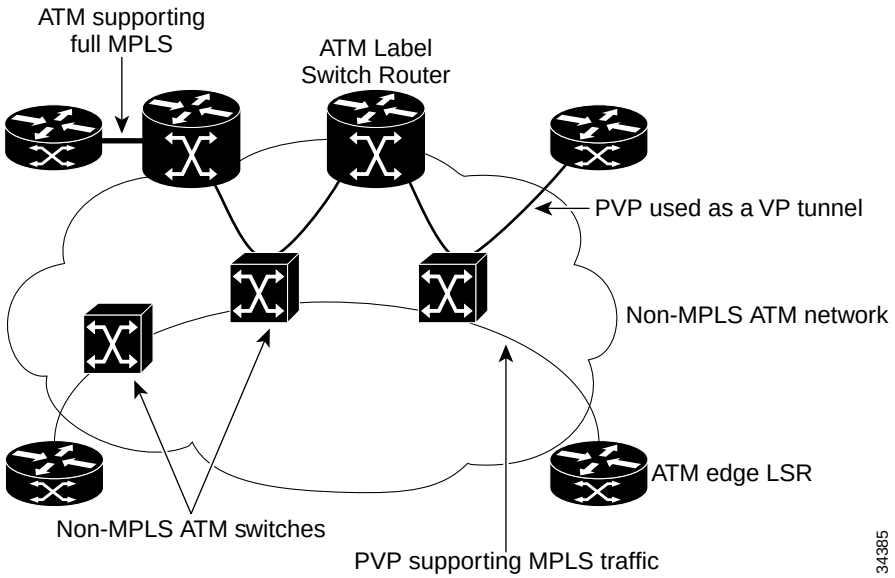

It is important to note that the Edge LSRs are part of the provider network and are controlled by the provider. The edge LSRs are critical to network operation and are not intended to be CPE under any circumstances. The provider may locate and manage routers at customer sites, but these are running ordinary IP and are outside the MPLS network.

MPLS Applications

MPLS networks as shown in Figure 1-1 have three main applications. Typically, two or all three of these capabilities would be used simultaneously:

- **IP+ATM Integration**
  MPLS fully integrates IP services directly on ATM switches. The IP routing and LDP software resides directly on ATM switches. Thus MPLS allows ATM switches to optimally support IP multicast, IP class of service, RSVP, and Virtual Private Networks (see below). Optimal integration of IP+ATM means that MPLS is far more scalable and far less complex than overlay schemes like MPOA, CSI, and IP Navigator.
- **IP Virtual Private Network (VPN) Services**
  A VPN service is the infrastructure of a managed Intranet or Extranet service offered by a provider to many corporate customers. These are often massive IP networks. MPLS, in combination with the Border Gateway Protocol (BGP), allows one provider network to support thousands of customer's VPNs. In this way, MPLS with BGP offers a very flexible, scalable, and manageable way of providing VPN services on both ATM and packet-based equipment. Even on small provider's networks, the flexibility and manageability of MPLS+BGP VPN services are a major benefit.
- **IP Explicit Routing and Traffic Engineering (TE)**
  An important problem in current IP networks is the lack of ability to finely adjust IP traffic flows to make best use of available network bandwidth. Also absent are related capabilities to send selected flows down selected paths, for example, to select protected trunks for particular classes of traffic. MPLS uses Label Switched Paths (LSPs), a type of lightweight VC. These can be set up on both ATM and packet-based equipment. The IP Traffic Engineering capability of MPLS uses special LSPs to finely adjust IP traffic flows.

The following summarizes label switching operations in various network services. More specific descriptions are covered in subsequent chapters.

MPLS Virtual Private Network

MPLS Virtual Private Networks (VPN) deliver enterprise-scale connectivity deployed on a shared infrastructure with the same policies enjoyed in a private network. A VPN can be built on the Internet or on a service provider's IP, Frame Relay, or ATM infrastructure. Businesses that run their intranets over a VPN service enjoy the same security, prioritization, reliability, and manageability as they do in their own private networks.

VPNs based on IP can extend intranets over wide-area links to remote offices, mobile users, and telecommuters. They can support extranets linking business partners, customers, and suppliers to provide better customer satisfaction and reduced manufacturing costs. VPNs can also connect communities of interest, providing a secure forum for common topics of discussion.

New IP-based services such as video conferencing, packet telephony, distance learning, and information-rich applications offer businesses the promise of improved productivity at reduced costs. As these networked applications become more prevalent, businesses increasingly look to their service providers for intelligent services based on a rich set of controls that go beyond transport to optimize the delivery of applications end to end. Today organizations want their applications to traverse a network in a secure, prioritized environment, and they want the opportunity to reduce costs, improve connectivity, and gain access to networking expertise.

Intranet and Extranet VPNs

Intranet VPN services link employees, telecommuters, mobile workers, remote offices, and so on, to each other with the same privacy as a private network.

Extranet VPN services link suppliers, partners, customers, or communities of interest over a shared infrastructure with the same policies as a private network.

Cisco provides a range of ATM- and IP-based choices for deploying large-scale intranet and extranet VPN services, including Multiprotocol Label Switching (MPLS)-based services, which provide secure, business-quality VPN solutions that scale to support tens of thousands of VPN customers over IP or IP+ATM networks.

A VPN built with MPLS affords broad scalability and flexibility across any IP, IP+ATM, or multivendor backbone. MPLS forwards packets using labels. The VPN identifier in the label isolates traffic to a specific VPN. In contrast with IP tunnel and virtual-circuit architectures, MPLS-based VPNs enable connectionless routing within each VPN community. Service providers can easily scale their services to support tens of thousands of VPNs on the same infrastructure, with full QoS benefits across IP and ATM environments.

Cisco MPLS-based VPN solutions are supported on its IP+ATM WAN switch platforms including the BPX 8650 and MGX families, and on its high-end router platforms such as the Cisco 12000 series GSR.

MPLS VPN Features

The VPN feature for MPLS Switching allows a Cisco IOS network to deploy scalable IPv4 Layer 3 VPN backbone services. MPLS Switching VPNs provide essential characteristics and features that service providers require to deploy scalable VPNs and build the foundation to deliver these value-added services:

Performance

When MPLS VPNs are set up using ATM LSRs such as the BPX 8650, the capabilities of scalable connectionless service of IP are combined with the performance and traffic management capabilities of ATM.

Connectionless Service

A significant technical advantage of MPLS VPNs is connectionless service. The Internet owes its success to its basic technology, TCP/IP, built on the packet-based, connectionless network paradigm. This means that no prior action is necessary to establish communication between hosts, making it easy for two parties to communicate.

To establish privacy in a connectionless IP environment, current VPN solutions impose a connection-oriented, point-to-point overlay on the network. Even if it runs over a connectionless network, today's VPN cannot take advantage of the ease of connectivity and multiple services available in connectionless networks.

By creating a connectionless MPLS VPN, tunnels and encryption are not required for network privacy, thus eliminating significant complexity.

Centralized Service

Building VPNs in Layer 3 has the additional advantage of allowing delivery of targeted services to a group of users represented by a VPN.

A VPN must give service providers more than a mechanism for privately connecting users to intranet services. It must also provide a way to flexibly deliver value-added services to targeted customers. Scalability is critical, because customers want to use services privately in their intranets and extranets.

Because MPLS Switching VPNs are seen as private intranets, it's easy to leverage new IP services:

- multicast
- Quality of Service
- telephony support within a VPN
- centralized services such as content and Web hosting to a VPN

Now myriad combinations of specialized services can be customized for individual customers, for example, a service that combines IP multicast with a low-latency service class to enable video conferencing within an intranet.

Scalability

Scalability is the major deficiency of VPNs created using connection-oriented, point-to-point overlays, Frame Relay, or ATM VCs. Specifically, connection-oriented VPNs require a full N^2 mesh of connections between customer sites to support any-to-any communication.

MPLS-based VPNs instead use the peer model and Layer 3 connectionless architecture to leverage a highly scalable VPN solution. The peer model requires a customer site to make peer connection with only one provider edge (PE) router as opposed to all other CPE or customer edge (CE) routers that are members of the VPN. The connectionless architecture allows the creation of VPNs in Layer 3, eliminating the need for tunnels or VCs.

Other scalability capabilities of MPLS Switching VPNs are due to the partitioning of VPN routes between PE routers and the further partitioning of VPN and IGP routes between PE routers and provider (P) routers in a core network. PE routers must maintain VPN routes for those VPNs who are members. P routers do not maintain any VPN routes. This increases the scalability of the providers core and insures that no one device is a scalability bottleneck.

Security

MPLS Switching VPNs offer the same level of security as connection-oriented VPNs. Packets from one VPN will not inadvertently go to another VPN. Security is provided at the edge and core of a provider network:

- at the edge, security ensures that packets received from a customer are placed on the correct VPN
- at the backbone, VPN traffic is kept separate

Malicious spoofing of a provider edge (PE) router is nearly impossible because the packets received from customers are IP packets. These IP packets must be received on a particular interface or subinterface to be uniquely identified with a VPN label.

Easy to Create

To take full advantage of VPNs, it must be easy to create new VPNs and user communities. Because MPLS VPNs are connectionless, no specific point-to-point connection maps or topologies are required.

Now it is easy to add sites to intranets and extranets and to easily form closed user groups. Managing VPNs in this manner enables membership of any given site in multiple VPNs, maximizing flexibility in building intranets and extranets.

Flexible Addressing

To make a VPN service more accessible, users should be able to design their own addressing plan, independent of addressing plans for other VPN customers supported by a common service provider.

Many organizations use private address spaces, as defined in RFC 1918 today, and do not want to undertake the time and expense of implementing registered IP addresses to enable intranet connectivity. MPLS VPNs allow customers to continue to use their present address spaces without network address translation (NAT) by providing a public and private view of the address.

If two VPNs want to communicate and both have overlapping addresses, that communication requires NAT at one endpoint. This enables customers to use their own unregistered private addresses and communicate freely across a public IP network.

Integrated Class of Service (CoS) Support

CoS is an essential ingredient of an IP VPN because it provides the ability to address two fundamental VPN requirements:

- predictable performance and policy implementation
- support for multiple classes of service in an MPLS Switching VPN

Network traffic is classified and labeled at the edge of the network before traffic is aggregated according to policies defined by subscribers and implemented by the provider and transported across the provider core. Traffic at the edge and core of the network can then be differentiated into different classes by drop probability or delay.

Straightforward Migration

For service providers to quickly deploy these VPN services, a straightforward migration path is required. MPLS VPNs are unique because they can be built over multiple network architectures, including IP, ATM, Frame Relay, and hybrid networks.

Migration for the end customer is also simplified because there is no requirement to support MPLS on the customer edge (CE) router and no modifications are required to a customer's intranet.

MPLS VPN Benefits

- A platform for rapid deployment of additional value-added IP services, including intranets, extranets, voice, multimedia, and network commerce
- Privacy and security equal to Layer-2 VPNs by limiting the distribution of a VPN's routes to only those routers that are members of the VPN
- Seamless integration with customer intranets
- Increased scalability over current VPN implementations, with thousands of sites per VPN and hundreds of thousands of VPNs per service provider
- IP Class of Service (CoS), with support for multiple classes of service and priorities within a VPN, as well as between VPNs
- Easy management of VPN membership and easy provisioning of new VPNs for rapid deployment
- Scalable any-to-any connectivity for extended intranets and extranets that encompass multiple businesses
- MPLS enables business IP services

  - VPNs with strong SLAs for QoS
  - privacy and QOS of ATM without tunneling or encryption
  - enabled by Cisco's unique combination of MPLS and open standards routing
- Lower operating costs
  - enables low-cost managed services to increase SP market share
  - increases profits though lower marginal cost for new services
  - network establishes VPN connectivity; no provisioning
  - build once/sell many; single routing image for all VPNs
- The first transport-independent VPN
  - universal VPN: one VPN, any access/transport: dial, xDSL, ATM, and so on
  - service delivery independent of transport/access technology
- Simpler to use
  - VPN managed by the service provider
  - transparent support for private IP addresses
  - multiple QoS service classes to implement business net policy
- Revenue and growth
  - revenue from today’s transport services, growth from IP
- Business IP services enabled by MPLS/IOS
  - MPLS brings IOS to service provider ATM networks
  - MPLS is the new industry standard for bringing IP and ATM together
- Seamless service delivery
  - wide breadth of services; circuit emulation to IP VPNs
  - single pipe; multiple services (any service, any port)
- lower cost of operation and competitive advantages
  - ROI, TTM, economies of a multiservice network

References

- The Cisco “IP+ATM Solutions” page at http://www.cisco.com/go/ipatm has links to press releases, brochures, white papers and other information. Use the links on the left-hand side of the page.
- The OSPF version 2 specification is http://www.ietf.org/rfc/rfc2328.txt
- The “IS-IS for Routing in TCP/IP and Dual Environments” specification is http://www.ietf.org/rfc/rfc1195.txt
- IETF documents on MPLS are at http://www.ietf.org/html.charters/mpls-charter.html. The most important documents are:
  - “MPLS Architecture” draft-ietf-mpls-arch-05.txt
  - “MPLS Label Stack Encodings” draft-ietf-mpls-label-encaps-04.txt
  - “MPLS using LDP and ATM VC Switching” draft-ietf-mpls-atm-02.txt

  - "LDP Specification" draft-ietf-mpls-ldp-05.txt
  - "MPLS Support of Differentiated Services by ATM LSRs and Frame Relay LSRs" draft-ietf-mpls-diff-ext-01.txt
- Other IETF documents on Differentiated Services are at http://www.ietf.org/html.charters/diffserv-charter.html
- The most important IETF documents on the Border Gateway Protocol are:
  - "A Border Gateway Protocol 4 (BGP-4)" http://www.ietf.org/rfc/rfc1771.txt
  - "Multiprotocol Extensions for BGP-4" http://www.ietf.org/rfc/rfc2283.txt
  - A further informational document shows how BGP can be used to support VPNs: "BGP/MPLS VPNs," RFC 2457, http://www.ietf.org/rfc/rfc2547.txt
- The following books on routing, MPLS and related topics are very useful:
  - Halabi, B., *Internet Routing Architectures*, Cisco Press, 1997.
  - Metz, C., *IP Switching Protocols and Architectures*, McGraw-Hill, 1999
  - Rekhter, et al., *Switching in IP Networks*, Morgan Kaufmann, 1998
- Useful magazine articles are:
  - Feldman, et al., "Evolution of Multiprotocol Label Switching," *IEEE Communications Magazine*, Vol. 36, No. 5, May 1998
  - Metz, C., "Ingredients for Better Routing: Read the Label," *IEEE Internet Computin*g, Sept/Oct. 1998
- Some archives on MPLS and related technologies are:
  - http://infonet.aist-nara.ac.jp/member/nori-d/mlr/
  - http://dcn.soongsil.ac.kr/~jinsuh/home-mpls.html