White Paper

Marconi

THIS COULD BE YOUR FINEST HOUR.

Building Scalable Service Provider IP Networks

Connection-Oriented Networking Solutions

July 2000

Table of Contents

INTRODUCTION

Service providers today are experiencing unprecedented demand for IP services. According to a report by Forrester Research, Inc., the entire U.S. Internet Services industry in 1996 amounted to $1.3 billion. By 1998, just the business segment of the market reached nearly $4 billion. Service providers have responded to this demand by building out nation- and world-spanning networks. Since 1996, some service providers have had to double their backbone bandwidth nearly every four months just to keep up. Looking ahead, according to Forrester Research projections, the market for IP services in the U.S. will reach $57 billion by 2003, rivaling the amount businesses will spend on long-distance calls.

Factors contributing to the demand for bandwidth include[1]:

- Broadband access is reaching the American residential market. By 2003, 15% of American homes will have purchased broadband access to the Internet.

- More and more business is being done on the Internet. Firms in every industry, indeed all organizations, are looking at the Internet as a way to improve business processes and/or reduce the cost of doing business with partners. The big three auto makers - GM, Ford and Daimler-Chrysler, for example - have announced they'll form an online market with each other and their suppliers to drive down costs. In many cases, completely new approaches to traditional business practices are being introduced (e.g., Amazon.com, eBay, Yahoo! just to name a few).

- The growth rate of traffic on the Internet is about 100% per year, a much higher growth rate than for traffic on other networks. If present growth trends continue, data traffic in the U. S. will overtake voice traffic around the year 2002, and will be dominated by the Internet[2] (some analysts believe this may have already occurred).

- In addition to the U.S., the number of Internet users around the world is growing rapidly as well. The Computer Industry Almanac[3] has reported that by the year 2002, 490 million people around the world will have Internet access, or 79.4 per 1,000 people worldwide. That ratio grows to 118 people per 1,000 by year-end 2005. And despite the fact that the U.S. has an overwhelming lead in Internet users - nearly 43 percent of the total 259 million worldwide - The U.S. will have only one-third of the total Internet users in 2002, and that number is expected to decline to 27 percent by the end of 2005.

- Much of the current explosion in demand is fuelled by applications that exploit the "best effort" nature of today's Internet. But there is growing demand for services that require a higher level of capability, specifically higher predictability from the Internet. Such services include: commercial Layer 3 VPNs, intranets, extranets, out-sourced firewalling and encryption and Voice Over IP (VoIP) services. As the Internet changes into a new public network that supports these new applications, growth in demand for these "higher level" services is expected to accelerate.

- Service Level Agreements (SLAs) written to meet Layer 2, Layer 3, and even Layer 4 parameters are being requested by customers to support new, emerging bandwidth-hungry applications. Once the Internet is able to support these requirements, demand for these applications is expected to increase dramatically.

---

[1] *Internet Services Hyper Growth*, Forrester Research, Inc., March 1999
[2] *The Size and Growth Rate of the Internet,* K.G. Coffman and Andrew Odlyzko, March 1998
[3] *Computer Industry Almanac, 8th Edition*, 1999

- Web servers and, increasingly, application servers are not being located on an organization's own site, but rather at locations with high bandwidth and good connectivity to the core of the Internet.

- The demands of bandwidth from Applications Service Providers alone is expected to more than quadruple the current level of bandwidth demand in order for services to meet the performance requirements of their customers.

In this hyper-growth environment, service providers must find a way to accommodate the dramatic growth in network traffic and the number of users. To do so in a cost-effective way, they must add management capabilities and higher predictability to their IP networks. Predictability is critical to optimizing network capacity and providing premium revenue-generating IP services. At the same time, both the capital costs of growing the network and operational costs of offering an ever-widening selection of services must be minimized. Further, the solutions that service providers choose today must have a clearly articulated migration plan to incorporate future opportunities and technologies.

In many large IP networks, service providers have deployed connection-oriented ATM network cores to optimize bandwidth utilization and increase IP service predictability via ATM traffic engineering and infrastructure resilience. Though this has worked well, and indeed has enabled the current explosive growth of the Internet, there is a desire to simplify network operations by reducing the number of control planes operating in the network (currently ATM and IP).

One emerging technology, Multi-Protocol Label Switching (MPLS), has been widely identified as a new tool to help service providers meet the often-conflicting challenges of increased predictability, growth in revenue, and cost reduction. As will be discussed in this white paper, it is MPLS's connection-oriented nature that provides an ability to increase IP service predictability, create differentiated IP services, and potentially reduce operation costs in IP-centric and multi-service networks.

To achieve all of this, MPLS combines a variety of functions from both IP and ATM. Specifically, MPLS adds enhancements to IP routing protocols to make them connection-oriented. In short, MPLS aims to provide a Connection-Oriented Link Layer (COLL) for IP that results in reliable and predictable forwarding of IP traffic, and that enables traffic engineering, congestion management, optimized end-to-end transmission recovery, and differentiated IP services. Further, MPLS, when augmented with a QoS framework such as that specified in the IETF's Differentiated Services model, may enable deterministic QoS in IP-centric networks.

MPLS is the natural evolution required for networks to support predictable and optimized IP services, particularly in next-generation, IP-centric networks. The connection-oriented nature of MPLS aims to help service providers meet unprecedented customer demand and their own revenue and profit goals.

Internet Demand and Service Provider Network Evolution

In the last ten years, the Internet has grown from a small network of interconnected routers to a world-spanning network that global businesses are coming to rely on as a mission-critical tool. Table 1 shows a high-level overview of the control plane protocols involved, the data plane data-transfer format involved, IP's strengths that led in part to the success of the early Internet, and the limits of IP in today's rapidly expanding and changing Internet.

Table 1: IP's Characteristics

	IP
Network control plane	
Admission control	None
Routing	OSPF, IS-IS, BGP4
Path computation	None—per hop forwarding
Signaling	None—per hop forwarding
Connection Name	None
Connection ID	None
Explicit Routing	None
Network data plane	
Transmission unit	Packets (variable length)
Policing (for fairness)	None
Marking	None
Buffer allocation	Limited
Scheduling (for flow prioritization and fairness)	Limited-none set by protocol standards
Strengths in an IP-Centric Network	Flexibility; Rich suite of data-service protocols, UNIX OS integration; Multi-vendor, standards based implementation
Limitations in an IP-Centric Network	Limited support for differentiated, predictable services; Connectionless hop-by-hop routing creates congestion (hyper-aggregation) and under-utilization of network resources.

In just the last few years, the pace of the Internet's growth has seriously strained the capabilities of the traditional routed infrastructure. The concern for service providers has quickly become: *scaling the network to meet the growing demand, while improving service availability and minimizing end-to-end latency and operational costs.*

Traditional Routing

Every technology has its advantages and disadvantages. While being connectionless brings a number of well-known benefits to IP - for example, scalability and overall network resiliency - it has some drawbacks, most notably:

1) A tendency towards "hyper aggregation" of data on certain links, which leads to congestion,

2) a limited ability to alleviate hyper-aggregation by, for example, distributing traffic load over all available resources, and

3) an inability to provide "toll quality" service levels across a network end-to-end.

All three limitations are due to IP's connectionless nature, whereby traffic is transported on a hop-by-hop basis, with routing decisions made at every node.

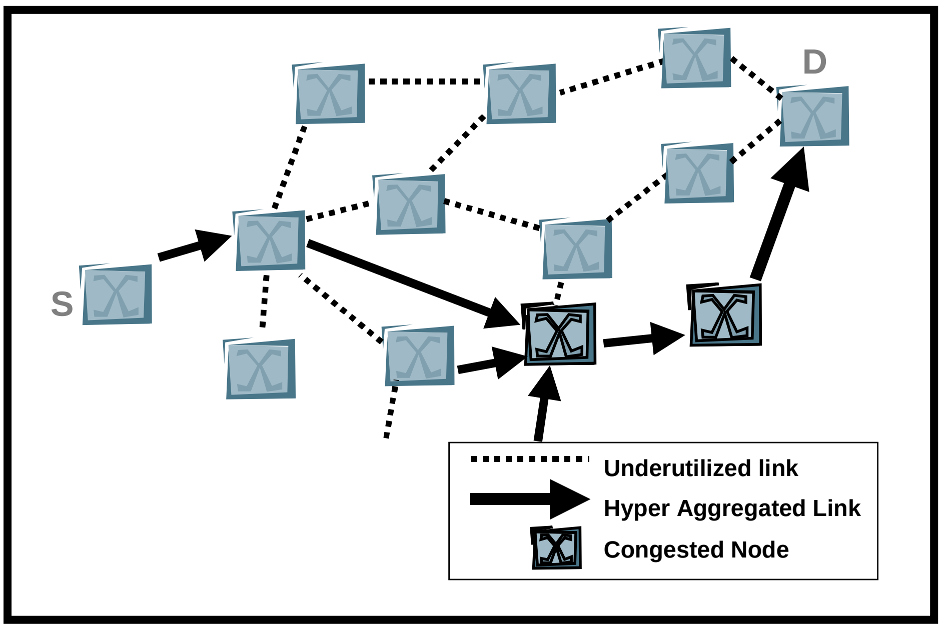

Figure 1: Hyper-Aggregation in an IP Network

Without a COLL, IP can create this kind of congestion in a network. Routers using the Open Shortest Path First (OSPF) routing protocol, for example, base routing decisions on the destination IP address of a packet's header, along with the least-cost path to that destination. All traffic then takes this least-cost path, congesting that path and leaving other paths through the network underutilized (see Figure 1). OSPF gives routers no end-to-end, overall view of the network and, therefore, the routers aren't aware of congestion in the network or of lightly loaded alternate routes and can't make the best use of all available network resources. Some of the larger ISPs claim that they lose up to forty percent of their network's capacity due to poor use of network resources by connectionless IP. By adding a COLL, they regain that capacity.

IP's hop-by-hop prioritization schema has another drawback: the inability to select paths through the network that guarantee the QoS requirements of latency-sensitive traffic flows. In an IP network, a real-time voice call or videoconference is routed by IP the same way as e-mail retrievals or bursty file transfers, so all three may experience congestion under conditions of hyper-aggregation.

While that may be fine for non-real-time data traffic like e-mail, a voice call or videoconference call has requirements for low-latency that must be met end-to-end across the network, from source to destination. IP routing protocols can't *guarantee* that these requirements will be met. Therefore, service providers running IP-centric networks can't make these guarantees on a network that scales over time in bandwidth, users, sites, and applications. They are limited to offering a “best effort” service. With no premium services to offer, IP service providers today are limited to charging only flat, commodity rates for the services they provide.

In addition, due to its connectionless nature, IP cannot guarantee fast per-flow reroute times for large-scale networks. Consequently, it’s difficult for IP service providers to offer any guarantee of network availability to customers looking for the “dialtone” service they’ve come to take for granted from the other large-scale network supporting their business: the telephone network.

In summary, IP is one of the most powerful networking technologies ever created, due to its adherence to open standards and its flexibility as a networking technology applicable to the transfer of a wide range of data types. However, in service provider networks, IP routing limits service providers’ ability to engineer and manage traffic in the network, and also limits the kind of service levels they can offer their customers. What’s required is a Connection-Oriented Link Layer (COLL) that is aware of the end-to-end state of the network, routes traffic based on the requirements of the application/user sending the traffic, and allows service providers to load-balance traffic across all available links in order to optimize the use of network resources. All of these would allow service providers to use a given network infrastructure as efficiently as possible, while making and *meeting* commitments to their customers. Today, that’s available with ATM and is emerging with MPLS.

The Benefits of a COLL include:

- Network devices that have knowledge of the network state, including the bandwidth available on each link.
- Devices that have knowledge of areas of congestion in the network.
- Devices that use end-to-end load balancing for optimal network bandwidth utilization.
- Virtual bandwidth partitioning that manages congestion and supports multi-tier service levels on common links,
- Optimized and predictable re-route times.

ATM In an IP-Centric Network

Many large-scale IP network operators have enhanced their IP service by incorporating ATM (Asynchronous Transfer Mode), a connection-oriented networking solution, in their networks. ATM optimizes network capacity through a Connection-Oriented Link Layer (COLL) that provides knowledge of the end-to-end state of the network in order to use the bandwidth on all available links optimally. Further, ATM’s connection-oriented nature enables virtual bandwidth partitioning to avoid congestion, and optimizes reroute times in the case of network failures. In addition, with current backbone interface speeds of OC-12c/STM-4 and OC-48c/STM-16, ATM has helped service providers meet the growing demand for Internet capacity.

Table 2 shows a high-level comparison of IP and ATM characteristics.

Table 2: IP and ATM Characteristics

	IP	**ATM**
<u>Network control plane</u>		
Admission control	None	UNI
Routing	OSPF, IS-IS, BGP4	PNNI
Path computation	None—per hop forwarding	End-to-End, constraint-based and congestion-aware
Signaling	None—per hop forwarding	PNNI
Connection Name	None	Virtual Connection
Connection ID	None	VPI, VCI
Explicit Routing	None	Designated Transit Lists
<u>Network data plane</u>		
Transmission unit	Packets (variable length)	Cells or packets (ATM Forum FAST)
Policing (for fairness)	None	Yes, for multiple traffic contracts
Marking	None	Cells are marked conform or non-conform
Buffer allocation	Limited	Per flow reservations
Scheduling (for flow prioritization and fairness)	Limited-none set by protocol standards	Per port, per flow, per class
Strengths in an IP-Centric Network	Flexibility; Rich suite of data-service protocols, UNIX OS integration; Multi-vendor, standards based implementation	Network predictability and reliability; Mature, field-hardened solutions; Connection Oriented, Layer2/Layer3 network partitioning; Optimized network bandwidth utilization; End-to-end load balancing
Limitations in an IP-Centric Network	Limited support for differentiated, predictable services; Connectionless hop-by-hop routing creates congestion (hyper-aggregation) and under-utilization of network resources.	Additional control plane to manage; Lack of ATM integration in routers results in a large number of router adjacencies to manage.

Figure 2 illustrates a frequent implementation of ATM in a service provider network. The core, or backbone, of many service provider networks is made up of an intelligent mesh of ATM switches. This Connection-Oriented Link Layer (COLL) core provides the primary transport for IP traffic. Surrounding the ATM core is a ring of IP routers. Behind those rings of core routers, there are likely several other types of networks. They may be the next tiers of the Internet or service provider network hierarchy. In some cases they may be PoP (Points of Presence) LANs, where subscribers access the network as well as Web-hosting content and any applications (e.g., ERP) hosted at that location.

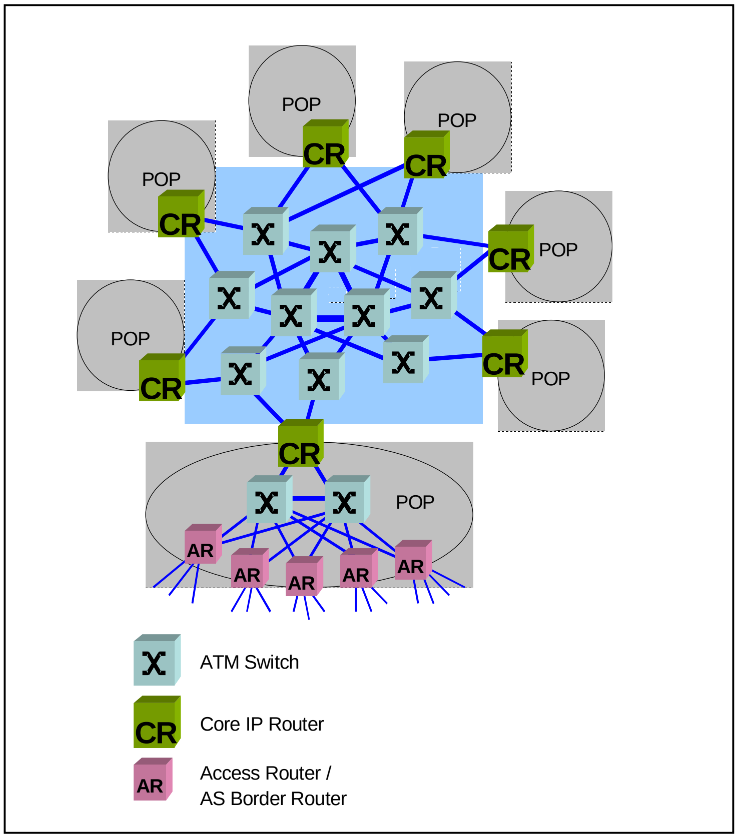

Figure 2: The Common Full-Mesh Configuration of Many Service Provider Networks

This network partitions Layer 2 and Layer 3 network functions, enabling the network to efficiently handle traffic at the networking layer best suited to the job. While the Layer 2 ATM switches provide the “big, fat pipes” to quickly and intelligently move data, the Layer 3 IP routers perform the IP routing and forwarding functions critical to the IP-centric network. In Figure 2, IP routers at the edge route data over the COLL provided by the ATM switches at the core. This offloads processor-intensive work from the routers, letting them simply map packets to the COLL of ATM virtual circuits - they now only have to deal with traffic as it enters or exits the core. In the traditional routed architecture, routers within the core also handle “transit” traffic. The amount of transit traffic increases with the size of the network, thereby stressing the routers. Offloading this traffic significantly enhances the scalability of the network.

Constraint-Based and Congestion-Aware Routing

Further, the ATM COLL provides two key functions that enable more predictable IP flows: minimizing congestion, and optimizing network capacity, or constraint-based routing and congestion-awareness. A device using constraint-based routing routes traffic based on traditional network topology information along with a number of other constraints, including the capacity and utilization of links, the requirements of the flow itself (i.e. bandwidth, delay, and jitter) and other administrative constraints. It may be used to guarantee specific applications (like video-conferencing) a fixed amount of bandwidth end-to-end through the network. It may also be used to minimize latency and jitter for voice traffic and to provide very specific, guaranteed and quantifiable customer service levels. This ability to configure varying ranges of quality of service to different customers of the network is also an attractive method for service providers to offer their customers differentiated services.

A congestion-aware device uses traditional routing information, but also takes into account the current state of the network with respect to traffic loading on each and every link. Congestion awareness provides network nodes the base information required to dynamically load-balance traffic through the network, and optimize traffic so as to avoid hyper-aggregation or under-utilization of links.

Constraint-based routing and congestion awareness are critical to ATM's connection-oriented performance. When data is sent across an ATM network, an end-to-end connection carries that data and takes into account:

1. The state of the network (available links, bandwidth available, guarantees available, etc.) and the most efficient routes through it.

2. The latency and bandwidth requirements of the application (or user) sending the data.

3. Preferred routes that have been configured previously by the network manager.

Service Availability and Network Resiliency

One of the ways service providers implement an ATM core is using Switched Permanent Virtual Circuits (SPVCs) between switches edge to edge. Using SPVCs, the network administrator simply designates the beginning and ending point of an ATM connection, while the rest of the connection in between is set up by the switches themselves, thanks to the high level of data each switch has about the network via PNNI.

Through PNNI, each switch has an extensive view of the network. Not only is it aware of the entire network topology, including the links, but also the speeds of those links, the current utilization of each, their delay, and other parameters that PNNI uses in routing updates. After performing sophisticated route computations with this information, PNNI selects an entire end-to-end path, a connection, which it encodes (by way of a Designated Transit List) into the signaling message. It then forwards application data out of the switch interface corresponding to the first switch entry on the path (this call set-up time in some vendor implementations is on the order of milliseconds. Once a call is set up, data transits a switch in microseconds). Any subsequent packets in that communication follow the same connection through the network.

SPVCs minimize operator involvement in configuring the network core and in provisioning services (in some vendor implementations, service provider customers themselves can determine and prioritize the switched routes used). In the event of a node or link failure, SPVC connections are automatically re-routed by the switches in milliseconds in the order of the connection priority set up by the administrator. Because the switches handle this rerouting themselves, it's very fast--on the order in some vendor's implementations of milliseconds. So fast that the ATM backbone of the service provider network has automatically detected a link failure and rerouted around it before the edge IP routers are even aware that there is a problem. This provides stability and reliability to the IP network that service provider customers can count on. Thus, ATM SPVCs, fast call set-up, and connection prioritization enable service providers to easily provision their network, offer high network availability to their customers and, potentially, offer premium-priced differentiated services.

In addition, because the ATM COLL provides congestion awareness, SPVCs also enable load-balancing of traffic across multiple links in the network. Rather than try to send a number of large data transfers down a congested path, for example, SPVCs allow the network manager to balance multiple flows across various links in the network and transfer the information without creating congestion.

Since the flows are switched via various end-to-end connections, some vendors optimize fairness between flows via per virtual connection (per VC) queuing and scheduling. An ATM virtual connection is the end-to-end connection over which an application's data flows through the network. By queuing and scheduling traffic into and out of device buffers on a per-VC basis, each VC is treated by the network the way it needs to be. So, for example, data generated by a videoconference is given high priority through device queues in order to provide smooth streaming images for the videoconference attendees. Per-VC queuing and scheduling can minimize delay for small packet flows at network congestion points, making sure delay-sensitive applications get treated in a way that optimizes their performance.

With per-VC queuing and buffering, it is also possible to partition network capacity and link bandwidth via connections that have specific, user-provided constraints. This partitioning allows operators to manage network traffic, protecting some high priority flows from the burstiness and congestion-creating volume of other flows.

This does not necessarily require policing of flows, but again is based on vendor implementation of per-VC queuing and prioritized scheduling.

A Limitation of IP Over ATM: Management of Multiple Control Planes

In an IP-centric network, network engineers and managers must configure, provision, and manage at least an IP topology. IP-over-ATM networks have an additional control plane to manage: the ATM control plane. This requires network managers to manage, provision, and control an ATM infrastructure in addition to their IP topology.

For some IP-centric operators, this control plane separation satisfies their desire to manage the backbone and infrastructure of the network separately from the IP service and access network, providing an additional level of reliability and stability to a critical part of the network (the backbone). For other IP-centric operators, the effort to manage this additional control plane either is not required. Perhaps highly predictable traffic flow is not required for the service profile being offered -- best-effort Internet access, for example, or is beyond the perceived benefits of the COLL. As we'll see in a moment, for these providers, MPLS has the potential to solve this problem.

Another limitation for some is the much publicized "cell tax" created when breaking down IP packets into 53-byte cells comprised of 48 byte payloads and 5 bytes of overhead. A new, open standard solution is "frame ATM", which uses variable-length ATM frames rather than fixed-length cells to transmit data. In the case of frame ATM, the cell tax is effectively "repealed." This is helpful for IP-centric providers who trunk traffic at OC-3c or OC-12c rates.

On the other hand, some operators are not concerned with this overhead for one of two reasons: either the traffic management advantages of ATM (including capturing underutilized network bandwidth) outweigh the additional overhead or, for the access portion of multi-service networks (those aggregating voice, video and data over the same links, for example, at speeds below 622 Mbps), they find the fixed 53-byte size of a traditional ATM cell is required to ensure all traffic receives the service it requires from the network, particularly in terms of delay.

MPLS: CONNECTION-ORIENTED NETWORKS

As a connection-oriented technology, MPLS:

- Enables IP-centric networks to be more predictable and efficient through load-sharing of traffic across multiple links in a network, and by using a COLL that enables network resources to be used more efficiently,
- Enables network managers to avoid hyper-aggregation scenarios by providing some level of traffic engineering and traffic management,
- Enables service providers to offer higher levels of service to their customers by allowing high priority flows in the network to receive higher prioritization than others,
- May enable service providers to offer high network availability if the network elements can provide fast rerouting of Label Switched Paths, and
- Reduces the number of control planes to be managed in an IP-over-ATM network.

MPLS Makes IP-Centric Networks More Predictable and Reliable

MPLS introduces more predictability and reliability to IP-centric networks through traffic engineering and traffic management functionality enabled by MPLS's standards-based extensions to IP routing protocols. Those extensions effectively provide IP traffic the benefits of a COLL. The MPLS control plane sets up MPLS connections - known as Label Switched Paths (LSPs) - from ingress Label Switching Routers (LSRs) at the network edge through the core LSRs to form a connection across a service provider's network.

MPLS running on IP routers enables network managers to assign traffic to LSPs based on information about the end-to-end state of the network. That avoids the hyper-aggregation and under-utilization of IP routing by directing traffic away from congestion-prone links onto less-congested links. This allows alternate paths to be configured, either manually or automatically, across both IP and ATM portions of the network for fast rerouting in the case of node failure, providing high network availability.

> The IETF is working on new standards that will enable IP and MPLS to become the common control plane for other transport technologies, such as DWDM. In this case, lambdas would be treated as LSPs controlled by MPLS.

In addition, as the MPLS standard matures and field-tested solutions emerge, MPLS may enable applications like voice and real-time video to get the network resources they require. An MPLS Forwarding Equivalence Class (FEC) enables customer traffic to be mapped to Label Switched Paths for high class of service traffic, thus enabling service provider networks to meet the latency and delay-tolerance requirements of these delay-sensitive applications. Thus, MPLS could enable service providers to make "high-margin promises" in a Layer 2 or Layer 3 Service Level Agreement (SLA), then confidently *meet* those promises in real-world networks.

MPLS Provides a Single Control Plane to Manage

When IP routers and ATM switches are both running MPLS, both the edge and core of the network may operate using the MPLS control plane. MPLS Label Switched Paths (LSPs) that pass through the entire network - from IP routers through ATM switches and back into routers - can be set up using only MPLS routing protocols. No longer will edge routers be running OSPF or IS-IS while the core ATM switches run PNNI. Now both sets of devices will operate using MPLS routing protocols so service providers have only one control plane to manage.

MPLS Signaling Protocols: RSVP, RSVP-TE, LDP and CR-LDP

As MPLS becomes a more robust standard and products emerge that run MPLS, different networking solution vendors will offer slightly different versions of MPLS. As customers begin to choose among these different MPLS offerings, one item of interest will be which signaling protocols a vendor supports.

> Another advantage of MPLS is that it minimizes the IP lookup, forwarding, and classification process. Rather than each router performing this function, it needs to be done only once-at the ingress to the MPLS connection and once again at the egress of the connection.
>
> Couple this with the actual connection (LSP) that provides a traffic-engineered path which guarantees the necessary bandwidth, and MPLS becomes a very exciting technology that marries the best of IP routing and ATM switching.

The two primary options are RSVP (Resource Reservation Protocol) and LDP (Label Distribution Protocol). When choosing an MPLS vendor, it is especially important to select one that has implemented the open-standards-based MPLS traffic engineering, traffic management, Quality of Service, and Constraint-Based Routing extensions not only to the IP routing protocols, but also to the signaling protocol.

An exciting possibility for MPLS signaling is its use to support the IETF IntServ (Integrated Services) QoS (Quality of Service) standard. Rather than simply using RSVP-TE or LDP to signal and maintain MPLS flows, these signaling protocols could be used for QoS management of traffic on aggregated trunks, enabling the MPLS COLL to provide QoS for IP traffic. This work is fairly new, but looks to become quite important. Many vendors, including Marconi, plan to integrate IntServ QoS and DiffServ with their RSVP-TE and LDP implementations.

Table 3 gives a high-level summary of the characteristics of IP, MPLS and ATM.

Table 3: IP, MPLS and ATM Characteristics

	IP	MPLS	ATM
Network Control Plane			
Admission control	None	*Not yet set forth in the MPLS standard	UNI
Routing	OSPF, IS-IS, BGP4	OSPF-TE, IS-IS-TE, BGP4-TE	PNNI
Path computation	None—per hop forwarding	End-to-End, constraint-based and congestion-aware	End-to-End, constraint-based and congestion-aware
Signaling	None—per hop forwarding	RSVP-TE, CR-LDP	PNNI
Connection Name	None	Label Switched Path	Virtual Connection
Connection ID	None	Label ID	VPI, VCI
Explicit Routing	None	Explicit Route Objects	Designated Transit Lists
Network data plane			
Transmission unit	Packets (variable length)	Packets and/or Cells	Cells or packets (ATM Forum FAST)
Policing (for fairness)	None	None	Yes, for multiple traffic contracts
Marking	None	None	Cells are marked conform or non-conform
Buffer allocation	Limited	*Not yet set forth in the MPLS standard	Per flow reservations
Scheduling (for flow prioritization and fairness)	Limited-none set by protocol standards	*Not yet set forth in the MPLS standard	Per port, per flow, per class
Strengths in an IP-Centric Network	Flexibility; Rich suite of data-service protocols, UNIX OS integration; Multi-vendor, standards based implementation	Efficient provisioning; Predictability and reliability; Support for differentiated services and SLAs; Optimized network bandwidth utilization; End-to-end load balancing	Network predictability and reliability; Mature, field-hardened solutions; Connection Oriented, Layer2/Layer3 network partitioning; Optimized network bandwidth utilization; End-to-end load balancing
Limitations in an IP-Centric Network	Limited support for differentiated services; Limited predictability, Less optimized networks; Connectionless hop-by-hop routing creates congestion (hyper-aggregation) and under-utilization of network resources.	An emerging standard; Little field experience; Lack of policing minimizes ability to guarantee fairness and throughput; No current plans for multi-service support or interoperability, therefore cannot currently be proposed as ubiquitous, common backbone for multi-service network operators.	Additional control plane to manage; Lack of ATM integration in routers results in a large number of router adjacencies to manage.

*The MPLS Forum is currently developing documents to aid service providers in determining the strength of individual vendor's implementations of these features. These documents will serve a similar function to the ATM Forum PICS documents.

Operating an MPLS Network

Like ATM VCs before them, MPLS LSPs allow service providers to increase control over their networks and, potentially, also offer various levels of service (e.g., gold, silver, bronze) to customers. LSPs can be set up by the network administrator, for example, to reduce the chances of hyper-aggregation in the network. On the other hand, since LSRs are aware of bandwidth on all links in the network, LSRs can set up LSPs that route traffic in ways that make most efficient use of all the links in the network, avoiding congestion entirely. In addition, LSRs can set up paths that are constrained by various user or application requirements. For example, if a user has purchased a “gold” level of service, high-priority LSPs could be set up for their data traffic to ensure that they always get the gold level of service from the network. In this way, MPLS may enable the incorporation of the IETF’s DiffServ standard in service provider networks.

LSPs can be set up to operate in several different ways:

- Point-to-Point switched paths, for example, can be used to connect all ingress nodes to all egress nodes to transport unicast traffic.

- A multipoint-to-Point LSP can connect all ingress nodes to a single egress node. This allows many microflows to take the same path through the network when they’ve been assigned to one Forwarding Equivalence Class (FEC).

- A multipoint-to-multipoint switched path can be used to combine multicast traffic from multiple sources into a single multicast distribution tree through the network.

- LSPs can also be tunnels, using label switching rather than network-layer encapsulation (L2TP, PPTP) as the means of moving packets through the tunnel.

In addition, MPLS allows streams of data to be forwarded as a unit, along a LSP. Thus a LSP through the network can be a single flow of user data or an aggregate flow of many users’ data. MPLS uses the term “Forwarding Equivalence Class” or FEC to refer to a set of Layer-3 packets that are forwarded in the same manner by a particular MPLS node. The mapping of IP packet to a FEC occurs only once per LSP, at the ingress LSR of the path. The LSRs in the core of the network simply switch on the MPLS header already applied by the ingress LSR.

MPLS is also a possible mechanism for provisioning VPNs. Packets coming from a customer network, for example, could contain an encapsulated header with a VPN label. At the service provider VPN ingress node, the header could be removed and a VPN label applied for switching through the VPN. At the VPN egress, the VPN label would be removed and the original label would provide label switching through the customer site.

Misunderstandings About MPLS

One misunderstanding about MPLS is that it is “just IP”. Another misunderstanding is that it is “just like ATM” and the two are mutually exclusive and/or that running both in a network is redundant. Neither statement could be further from the truth. Instead, MPLS is a Connection-Oriented Link Layer (COLL) over which IP can run. As such, it brings some of the benefits of connection-orientation to an IP-centric network. MPLS is a kind of “common ground” between IP and ATM (an OSI “Layer 2.5”, if you will). It replaces neither IP nor ATM, but is a new networking tool that optimally solves a certain set of network problems. It combines functions of IP with functions of ATM to provide a new tool for 21st century networking solutions.

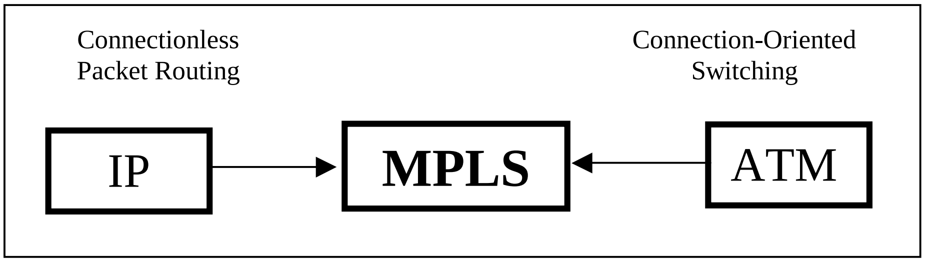

Figure 3: MPLS—A New Complimentary Networking Tool

MPLS, however, is a new standard and not yet field-tested. As such, it provides some level of COLL benefits today to IP services and networks, with more promised through future standardization. Some of those benefits include:

- In an all Packet Over SONET data network, MPLS can provide connections to enable a level of traffic engineering and traffic management that IP alone cannot provide. This is the type of network to which MPLS can add the most value.
- MPLS operates over frames or cells so it can operate in IP-over-ATM networks as well as IP-over-MPLS networks. This “protocol agnosticism” enables service providers to continue to run the IP over ATM networks they have today and slowly add MPLS to parts of the network as their business plans require. Additionally, this can enable a step-by-step approach to migrating an IP over ATM network to all MPLS without entailing the risk of a hard cut-over from one type of network to the other.
- MPLS has some congestion awareness and its routing protocols (IP protocols with traffic engineering extensions) are constraint-based. Hence, it should be able to provide some level of traffic engineering and management that reduces hyper-aggregation and should be able to provide some level of Class of Service prioritization to traffic. When combined with the IETF’s Differentiated Services model, it may also enable service providers to guarantee a higher level of service to customers who require it.
- In an IP over ATM network, MPLS reduces the two control planes needed to run the network-IP and ATM- to one control plane, that of MPLS. This simplification of network operation promises lower operational costs for service providers, enabling them to achieve higher profit levels from the network infrastructures they’ve already put in place.

Additionally, there is room left for MPLS to grow to provide all the benefits of a COLL. For example:

- Unlike ATM, which has quantitative, “hard” QoS capabilities built into standards, as well as built into vendors’ implementations, MPLS currently holds only the promise of providing IP traffic some level of “soft” QoS. Likely this will at first be some level of high or low prioritization given on a qualitative Class of Service (CoS) basis.

- Both ATM and MPLS are sensitive to calls per second setup performance for optimizing service availability during re-routes around failures. Best-of-breed vendor ATM implementations, for example, can perform a reroute in less than 50 milliseconds. It remains to be seen how quickly MPLS Label Switched Paths (LSPs) can be rerouted.

- The MPLS standards have not specified standardized COS/QOS parameters, as has already been done with ATM. As such, multi-vendor interoperability for MPLS-based circuits (LSPs) providing COS or QOS may not be available in the first phase of MPLS network roll-outs.

MPLS is a new standard. Vendor implementations on routers will not provide all the traffic management benefits seen on some vendors' ATM or MPLS switches. These benefits include traffic policing, traffic shaping, and hierarchical scheduling, which enable network operators to further engineer, manage, and control traffic in their networks.

Because of this, it's important to remember, that all vendors' MPLS implementations will not be equal. Taking advantage of the benefits of the COLL provided by MPLS requires more than just the implementation of the MPLS routing protocol. It also requires hardware that can provide fast setup of connections (MPLS calls-per-second), per Label Switched Path and micro-flow queuing, and per LSP and micro-flow scheduling. Only vendor offerings that implement this level of MPLS COLL can offer all the advantages that MPLS promises.

Integrating MPLS Into an IP Over ATM Network

MPLS offers the promise of enormous competitive advantages to service providers. Integrating MPLS into existing networks, however, should be done with care and, if possible, in a smooth, incremental way. Using the industry-standard "Ships In the Night" mode of operation can make such a migration from today's IP over ATM networks to MPLS a smooth one. Ships In the Night also allows service providers to provide non-IP transit services such as private line, frame relay, voice and ATM over the same physical networking infrastructure.

Ships in the Night operation allows a step-by-step migration between IP over ATM and MPLS without requiring a build-out of a parallel physical network. It allows service providers to create two logical networks on top of one physical topology by allowing IP over ATM and MPLS to run concurrently on the same devices. This is a very powerful configuration option that permits the same physical port to be configured with both IP/ATM **and** MPLS control planes **at the same time**. The multiple control plane protocols are able to function over a single physical network and be completely aware of each other.

This means that service providers will be able to add MPLS incrementally to their existing networks and not have to build a completely separate network for it. So rather than having to go to MPLS all at once, a service provider can simply "turn on" MPLS in small portions of the network as test beds, see how operations are effected, and move on from there in a methodical, controlled way, while customer traffic runs through the network.

In addition, there are many network designs where the backbone will need to support both ATM and MPLS simultaneously for an indefinite amount of time. This may be due to multiple services sold or provisioned, or due to varying technology support among multiple network edge devices (especially where multiple vendors are concerned). Ships in the Night operation in this situation presents a compelling solution.

CONCLUSION

In summary, many IP service providers have already chosen an IP-over-ATM network architecture to scale to meet their customers' unprecedented demands for bandwidth, to enjoy a high level of control over data traffic in the network through ATM's traffic engineering and traffic management capabilities, and to offer value-added, high-margin services to their Internet and WAN-access customers.

MPLS is a new technology that promises to achieve some level of this control in IP-only networks, and to simplify control plane management in an IP-over-ATM network. In an IP-only network, MPLS has the potential to reduce or remove the problem of hyper-aggregation of traffic on certain links, as well as offer more predictability and allow service providers to offer higher levels of service to their customers than traditional best-effort IP routing.

Additionally, for those service providers with IP-Over-ATM networks, and for service providers who wish to operate multi-service networks, the Ships in the Night mode offers an optimized multiservice network and a smooth migration path for IP-centric networks. Ships-In-the-Night allows those service providers to include MPLS capability in their network in a step-wise fashion or, should they choose to do so, migrate gracefully over time to an entirely MPLS-based network as the MPLS standard matures and as their own business model dictates.

IETF REFERENCES

Multiprotocol Label Switching Architecture, Internet draft draft-ietf-arch-04

A Framework for Multiprotocol Label Switching, Internet draft draft-ietf-mpls-framework-02.txt

MPLS Label Stack Encoding, Internet draft draft-ietf-mpls-label-encaps-03.txt

Requirements for Traffic Engineering Over MPLS, RFC 2702, draft-awduche-mpls-traffic-eng-00.txt

Extensions to RSVP for LSP Tunnels, Internet draft draft-ietf-mpls-rsvp-lsp-tunnel-01.txt

Multi-Protocol Lambda Switching, draft-awduche-mpls-te-optical-01.txt

Copyright © 2000 Marconi. All rights reserved. Marconi, FORE, AVA, CellPath, ForeRunner, ForeRunnerLE, ForeThought, ForeView and PowerHub are registered trademarks of Marconi and its affiliates. All Roads Lead To ATM, Application Aware, ASN, ASX, ATV, CellChain, CellStarter, Demarc, EdgeRunner, ERP Express, ESX, FramePlus, ForeRunnerHE, Intelligent Infrastructure, I^2, MSC, Netpro, Networks of Steel, Network of Steel, StreamRunner,TACtics Online,TNX,Universal Port and VoicePlus are trademarks of Marconi and its affiliates. All other brands or product names are trademarks of their respective holders. Please be advised that technical specifications set forth in this product literature are correct as of July 2000 and are subject to change by Marconi at any time without notice.