MPLS: Panacea Or Passing Fancy?

Bob Bellman

Is MPLS just the next big thing? Or does it offer real value and sustainable benefits?

In wide-area networking, as in many endeavors, less is more. If service providers could build complete wide-area networks (WANs) using only IP, their capital costs and, more importantly, their operational expenses would plummet. But IP isn't there yet. In particular, IP lacks the traffic engineering features that service providers need to sign service level agreements (SLAs) and still get a good night's sleep.

Multi-Protocol Label Switching (MPLS) promises to make all-IP networking a reality. By plotting static paths through an IP network, MPLS gives service providers the traffic engineering capability they require while also building a natural foundation for virtual private networks (VPNs). In addition, MPLS has the potential to unite IP and optical switching under one route-provisioning umbrella.

If you've been in the business for a few years, you probably remember similar world domination claims for ATM. And while ATM has gained popularity among service providers, it has not and will not achieve universal protocol status. Will MPLS—or more precisely, IP-over-MPLS—do any better?

From IP Switching To Traffic Engineering

Before considering where MPLS is going, it helps to remember where it came from. As a reaction to the IP switching threat from Ipsilon Networks back in 1996, virtually every major networking company rolled out a competing architecture. Cisco Systems announced Tag Switching, IBM published Aggregate Route-Based IP Switching (ARIS), Cascade Communications came up with IPNavigator, and so on.

All of these architectures were designed to make IP go faster, and while the details varied, all took the same basic approach: Use a standard routing protocol like Open Shortest Path First (OSPF) to lay out paths between endpoints; assign packets to these paths as they enter the network; and use ATM switches to move packets along the paths. ATM switches were much faster than IP routers back then, which was the whole point. By early 1997, the IETF proposed to define a standard version of IP switching and called it Multi-Protocol Label Switching. Working groups completed the core specifications for MPLS in the fall of 2000, and approved RFCs should be available by the time this article is published (see www.ietf.org).

FIGURE 1 IP-Over-ATM Requires Separate Virtual Circuits Among All Routers

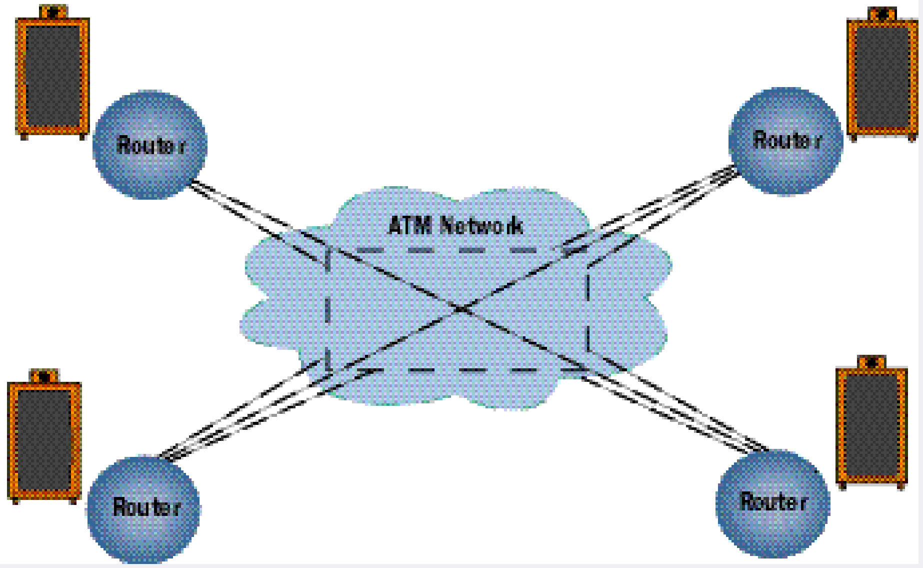

Bob Bellman is president of Brook Trail Research (www.brooktrail.com), a consultancy specializing in network technology assessment and marketing.

But the market didn't wait for MPLS. Between 1997 and 2000, eager engineers developed routers like the Juniper M40 and the Cisco 12000 that ran as fast as ATM switches. At the same time, someone noticed that MPLS was a convenient framework for both IP traffic engineering and VPNs. So MPLS lost its original *raison d'être*, but it gained two more sustainable purposes instead.

Less Than Perfect

The value of traffic engineering cannot be overstated, especially in the SLA-driven world of public networking. Simply put, traffic engineering allows service providers to do two things: control quality of service (QOS) and optimize network resource utilization.

As a connectionless protocol, IP does neither of these. Although it supports prioritization of packet streams, IP cannot guarantee that network resources will be available when needed. Moreover, IP sends all traffic between the same two points over the same route. During busy periods, therefore, some routes get congested while others remain underutilized. Without explicit control over route assignments, the provider has no way to steer excess traffic over less busy routes.

That's another reason so many service providers turned to ATM to transport their IP traffic: They needed ATM's virtual circuits to control bandwidth allocation on busy backbone routes. Because ATM is connection-oriented, it gives service providers the traffic engineering tools they need to manage both QOS and utilization. When provisioning an ATM network, the service provider can assign each virtual circuit a specific amount of bandwidth and a set of QOS parameters. The provider can also dictate what path each VC takes. Basing these decisions on overall traffic trends reduces the likelihood of network hot spots and wasted bandwidth. ATM's constraint-based PNNI routing also helps by automating the process. With IP alone, service providers stay up at night worrying about cost and performance. With IP-over-ATM, service providers have at least a chance of getting some sleep.

But IP-over-ATM is less than perfect. For one thing, the overlay approach means a network operator has to deal with two control planes, managing both IP routers and ATM switches. In the face of escalating operations costs, two control planes are one too many.

Moreover, using ATM VCs to interconnect IP routers leads to scaling problems, since every router needs a separate VC to every other router (Figure 1). As the network grows, the number of routes and VCs can increase exponentially, eventually exceeding the capacity of both switches and routers. Network operators could work around this problem, but only by forgoing a full-mesh architecture. MPLS addresses this problem directly, by bringing ATM-like connections under the control of IP routing protocols.

FIGURE 2 Label Switch Router

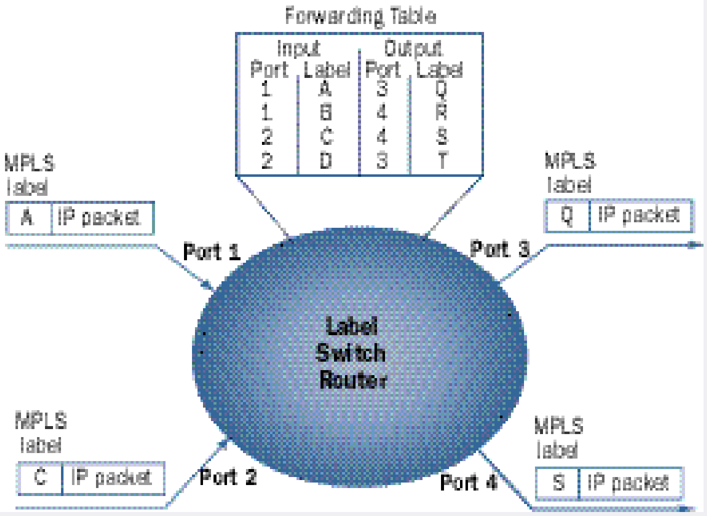

MPLS 101

An MPLS network comprises a mesh of *label switch routers* (LSRs), which are MPLS-enabled routers and/or MPLS-enabled ATM switches. As each packet enters the network, an ingress LSR assigns it a label based on its destination, VPN membership, type-of-service bits, etc. Then, at each hop, an LSR uses the labels to index a forwarding table (Figure 2). The forwarding table assigns each packet a new label—to promote scaling, labels have only local significance—and direct the packet to an output port. As a result, all packets with the same label follow the same *label switched path* (LSP) through the network.

One important difference between MPLS and ordinary IP routing is that packets sent between the same two endpoints can have different labels, and thus can take different paths. Also, because LSRs look at a packet's label and no more, MPLS forwarding is simple and quick. Properly speaking, MPLS qualifies as a Layer-2 (or maybe 2.5) protocol, making it an IP-friendly successor to frame relay and ATM, and not a new generation of IP routing.

MPLS is IP-friendly because it uses the same routing protocols as IP, notably OSPF and IS-IS, to plot paths through the network. Moreover, in a network carrying both ordinary IP and MPLS traffic, LSRs and ordinary routers exchange route updates as peers. Of course, the routers consult their route tables each time they forward a packet, while the LSRs use their route tables only when provisioning LSPs. Still, the result is a unified control plane.

Traffic Engineering At Last

To support traffic engineering, MPLS lets service providers specify explicit routes for LSPs. Using explicit routes, service providers can reserve network resources for high-priority or delay-sensitive flows, distribute traffic to prevent network hot

FIGURE 3 Label Stacking

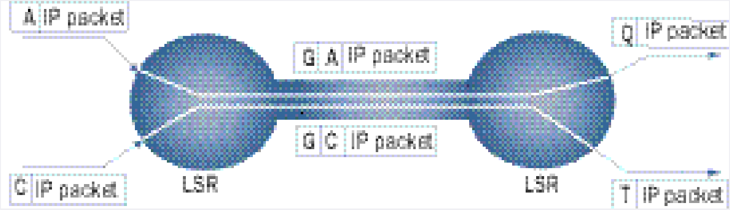

spots and pre-provision backup routes for quick recovery from outages. Explicit routing finally gives network operators the hooks they need to optimize bandwidth utilization and control IP QOS without resorting to ATM.

Network operators can specify explicit routes manually, by configuring them into edge LSRs, or they can use one of two new signaling protocols—RSVP-TE and CR-LDP—to automate the process. Either protocol can be used by the ingress LSR to tell the network what route a new LSP must follow, how much bandwidth to reserve for that path, and other QOS requirements.

As their names imply, RSVP-TE is conventional RSVP with traffic engineering extensions, while CR-LDP is LDP—the MPLS *label distribution protocol*—augmented for constraint-based routing. For a while, the two protocols were a bone of contention. The RSVP-TE camp argued that RSVP was a proven signaling protocol, so why invent something new? CR-LDP proponents maintained that LDP was designed specifically for MPLS, and anything else would be a compromise.

Cisco chose to implement RSVP-TE first, and most equipment vendors followed suit. But the debate is winding down anyway. Jarrod Siket, director of product planning at Marconi Communications and vice-chair of the MPLS Forum technical committee, explained the situation as follows: “RSVP-TE and CR-LDP are very similar. It’s absurd to argue that one is right or wrong, and we can’t tell a customer which to use anyway. Marconi will do both.” So will most other equipment vendors.

Although traffic engineering is essential to MPLS, it is not mandatory. Service providers don’t have to use explicit routing, and probably won’t in networks with plenty of bandwidth. Instead, they can let ingress LSRs use LDP—without any constraint-based extensions—to automatically associate labels with paths. With plain LDP, MPLS packets follow the same routes as ordinary routed packets.

Hop Popping

A less well-known feature of MPLS, *label stacking*, may turn out to be one of its most powerful attributes. Label stacking lets LSRs insert an additional label at the front of each labeled packet, creating an encapsulated tunnel that can be shared by multiple LSPs (Figure 3). At the end of the tunnel, another LSR pops the label stack, revealing the inner label. An optimization in which the next-to-last LSR peels off the outer label is known in IETF documents as “penultimate hop popping.” (Someone has been reading too much Dr. Seuss!)

Unlike ATM, which has only one level of stacking—virtual channels inside virtual paths—MPLS supports unlimited stacking. Thus an enterprise could use label stacking to aggregate multiple flows of its own traffic before handing it to an access provider; the access provider could aggregate traffic from multiple enterprises before handing it to a backbone provider; and the backbone provider could aggregate yet again before passing traffic to a wholesale carrier.

Service providers could use label stacking to merge hundreds of thousands of LSPs into a relatively small number of backbone tunnels between points of presence. Fewer tunnels means smaller route tables, making it easier for providers to scale the network core.

Building Better VPNs

Traffic engineering may be the operational *sine qua non* of MPLS, but VPNs top the marketing list of MPLS features. Rob Redford, director of marketing at Cisco, said MPLS lets service providers create VPNs with the flexibility of IP but the QOS of ATM. Separate labels ensure privacy between VPNs without resorting to encryption. “Then,” said Redford, “they can up-sell existing customers by adding new services to the VPN. It’s the only way that service providers can move up the value chain.”

Indeed, VPNs are the first application of MPLS for many carriers. Global communications provider Equant announced its MPLS-based VPN service in September 2000. Equant had offered IP VPN service for nearly five years, but was facing scaling problems. “We used BGP to create VPNs,” explained Jon Floyd, IP marketing manager, “so we had to deploy boundary routers for each VPN customer around the world. We couldn’t grow due to the physical space required by the boundary routers.”

Adding MPLS software to its Cisco edge routers allowed Equant to share them among different VPN customers. “Now we’re managing a smaller number of boxes,” said Floyd. “Also, now we can use unregistered private addresses.” Previously, Equant had to grant its customers unique IP addresses or use NAT (Network Address Translation) to isolate unregistered addresses. Now, MPLS labels hide users’ IP addresses, making NAT unnecessary.

AT&T (www.att.com) and AT&T Canada (www.attcanada.com) also offer MPLS-based VPNs. Both companies have offered frame relay and ATM services for years, but those services required customers to lease separate permanent virtual circuits between every pair of sites. To provide more flexible connectivity over the same backbone, AT&T Canada added router blades with MPLS software to its Cisco ATM switches and introduced its Business IP service. Now a customer needs only a single frame relay or ATM PVC at each of its sites to create a VPN. Inside the network, MPLS LSPs replace the point-to-point frame relay PVCs and provide complete connectivity between endpoints.

"Previously, a 10-site VPN would require 45 PVCs," said Doug Westlund, marketing VP at AT&T Canada. "With Business IP, the same customer needs just 10 PVCs. This simplifies the customer's network design, leads to smaller route tables, and lets the customer deploy less-costly premises equipment.

"The customer will see cost savings," Westlund continued, "but that's not our selling proposition." Instead, AT&T Canada promotes the connectivity, ease of use and class-of-service capabilities of MPLS. Business IP includes four classes of service: Service Class for best-effort traffic, Business Class for higher priority data, Voice Class for delay-sensitive traffic, and Premier Class for near real-time applications.

In the U.S., AT&T created a similar service by "IP-enabling" its existing ATM and frame relay switches. "Our customers are used to a certain level of performance with frame relay and ATM," says Tim Halpin, AT&T's product director for frame relay and ATM services. "We wanted to deliver the same level of service for IP VPNs." Besides knocking down access costs, the new MPLS-based VPN service makes it easier to add new customer sites—just provision a PVC from the new site to the VPN—and new services. "For example," said Halpin, "if a customer wants Internet access through a network-based firewall, it doesn't need a separate PVC to the firewall from each site. All sites have access to all services via the VPN."

MPLS Cell Switches?

Besides VPNs, scaling and traffic engineering, MPLS is good for at least one more thing: injecting new life into the ATM switch business. Adding MPLS support to an existing ATM infrastructure lets a carrier offer IP services without the scaling problems of IP-over-ATM. And the carrier can continue to support non-IP traffic over the same backbone.

Switch manufacturers like Cisco and Marconi refer to this approach as ATM+IP. The "plus" is significant. Even on ATM switches, MPLS-based IPdoesn't run over ATM. Instead, a router blade is typically added to the switch. The router runs MPLS routing and signaling protocols and merely exploits the speed and QOS properties of cell switching. MPLS and ATM act like "ships in the night" sharing network resources but never interacting.

One can imagine an ATM switch carrying nothing but MPLS traffic, but ATM switch vendors reject this scenario. "We don't see ATM going away anytime soon," said Marconi's Siket. "There are many revenue-generating services that can't run on IP." Cisco's Redford frames it differently, "With ATM+IP, a CLEC can build one network, deliver frame and ATM, and add IP for future services." Either way, ATM switch vendors keep selling hardware, and service providers extend the useful lives of their ATM gear. Of course, the providers are still stuck with dual control planes.

Proponents next want to extend MPLS into optical networks

Go To The Light

Not satisfied with conquering just IP and ATM, MPLS enthusiasts have marched on to attack optical networking as well. Within the IETF, an "MP-*Lambda*-S" initiative is looking at how the MPLS control plane can be extended into optical switches. "Optical switches are inherently connection-oriented," pointed out Ross Callon, distinguished engineer at Juniper Networks and a prime mover in the IETF. "What should you use to set up connections—PNNI, SS7 or something else? MPLS will succeed because it's a good method and it's getting the attention it needs to succeed."

In the grandest scenario, a unified MPLS control plane would span routers, ATM switches, SONET cross-connects and DWDM-based optical switches. All of the devices would cooperate as peers when laying out LSPs, and a router or cell switch could tell an optical network to provision light waves on demand.

It's a good idea, but that sort of protocol hegemony seldom, if ever, succeeds. For one thing, service providers are reluctant to give each other route-level visibility into their networks, so MPLS peering will likely stop at provider boundaries. Worse yet, organizational barriers may preclude MPLS peering *within* certain carriers; it seems unlikely that the transmission group will let the data services group dictate the allocation of light waves. Still, MP-*Lambda*-S could turn out to be the optical control plane of choice. After all, most of the bits traveling over light waves will ride in IPpackets.

Crossing The Demarc

Speaking of hegemony, is MPLS just for carrier backbones, or will it spread into enterprise networks as well? Clearly the benefits of MPLS will show through to enterprises, as carriers offer improved VPN services and extend their SLAs to include more meaningful performance-related criteria. Companies with private wide-area networks might also consider MPLS for their own use. "Really big private nets," says Juniper's Ross Callon, "are just like carrier nets and need MPLS. But you don't need traffic engineering or VPNs on a campus, so you don't need MPLS there."

You may, however, encounter MPLS where the campus meets the wide area, as multiple traffic flows are label-aggregated for handoff to the service provider, or as local access providers offer

FIGURE 4 Supporting Different Classes Of Servce

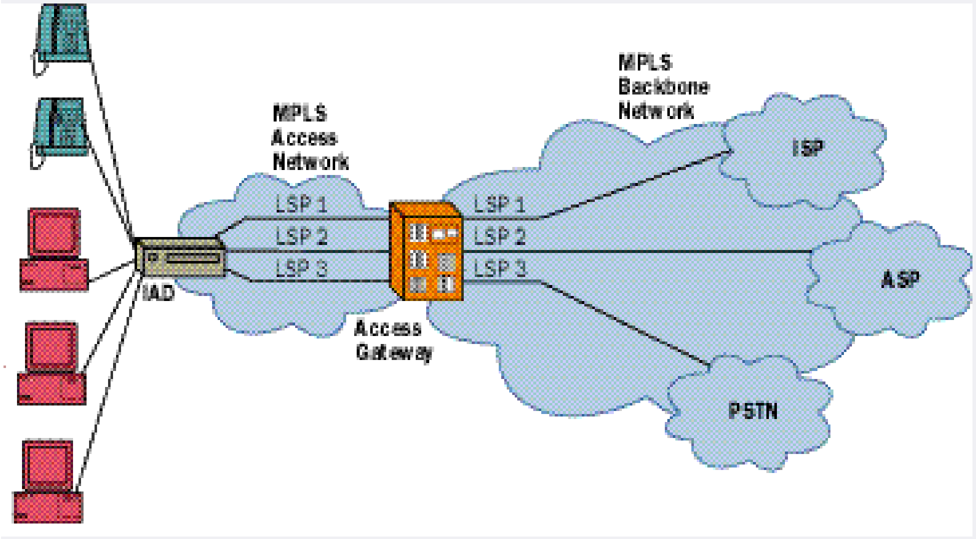

MPLS-based integrated access services. As shown in Figure 4, equipment-maker Integral Access, for example, relies on MPLS to carry IP-based voice and data between its integrated access device and its multiservice access gateway. The premises-based IAD classifies packets and sends them over separate LSPs for each class of service. The central office-based access gateway aggregates packets from multiple IADs and relays them to backbone LSPs.

MPLS makes it easy for Integral Access to separate voice from data and to give each flow the proper bandwidth and priority. "MPLS is also a nice way of isolating users as they go through the same switch," said Guy Chenard, VP of marketing and business development, "and MPLS is link-layer independent, so Integral Access can offer the same features over DSL, T1/E1 or fiber."

Encouraging Progress

Progress on the rollout of MPLS is encouraging. Dozens of service providers are experimenting with MPLS, and several have already put it into production. Multivendor interoperability tests have been consistently positive. Tests last fall at the Advanced Internet Laboratory (AIL) at George Mason University focused on basic signaling and traffic engineering. "ISPs are most concerned about traffic engineering, particularly recovery schemes," says Bijan Jabbari, founder and head of the AIL. The lab conducted a series of interoperability tests with equipment from Avici Systems, Cisco, Juniper and Nortel Networks. "The products worked very well together," reported Jabbari. "We achieved 100 percent interoperability by the end of the test."

Of course, MPLS is not yet fully baked. The IETF and MPLS Forum still have plenty of issues to deal with. For one, they must reconcile MPLS with DiffServ, so that type-of-service markings can be transferred from IP headers to MPLS labels and interpreted by LSRs in a standard manner. For another, they must clarify how MPLS supports virtual private networks. At least two models exist, one based on BGP and the other on virtual routers. The BGP model in particular is frighteningly arcane and could retard the implementation of MPLS-based VPNs by ordinary mortals.

Moreover, MPLS is not a seamless fit with today's networks, at least not for service providers. To realize the full benefit of MPLS, protocols like RSVP, OSPF and IS-IS must be extended, and network operators must master new concepts like explicit routing. They also have to work their way up to new business models. For AT&T Canada's new Business IP service, "We had to redefine the way we took orders," said Doug Westlund, "and we had to raise our knowledge of IP for the order-taking service." Presumably, only the benefits of these changes are visible to end users.

An Elegant Solution

Will MPLS take over the world? Definitely not. Nothing ever does, and IP visionaries are already promising new routing techniques that will make MPLS obsolete. Will MPLS replace ATM? Eventually, but only after MPLS-based QOS proves out. Does MPLS offer real value? Absolutely. MPLS is an elegant solution to a real-world problem—traffic engineering in IP networks—and it builds on lessons learned from years of IP, frame relay and ATM networking.

Without MPLS, service providers will continue to struggle with costly, inflexible network overlays. With MPLS, they can simplify their operational procedures, deliver more versatile IP services, sign meaningful SLAs and still get a good night's sleep□

Companies Mentioned In This Article

AT&T (www.att.com)
AT&T Canada (www.attcanada.com)
Avici Systems (www.avici.com)
Cisco Systems (www.cisco.com)
Equant (www.equant.com)
George Mason University (www.ail.gmu.edu)
IBM (www.ibm.com)
Integral Access (www.integralaccess.com)
Juniper Networks (www.junipernetworks.com)
Marconi www.marconi.com
MPLS Forum (www.mplsforum.org)
Nortel Networks (www.nortelnetworks.com)