Beyond The Core: MPLS In Metro, Access And Private Networks

Bob Bellman

Carriers will push the technology closer to the customer, but enterprises likely won't opt for an MPLS-UNI.

The battle for the backbone is over. Multiprotocol Label Switching (MPLS) has conquered the core of most public IP networks. Its label-switched paths (LSPs) are a natural fit for virtual private networks (VPNs). Its quality of service (QOS) and fast-reroute features let service providers sign service-level agreements with confidence. Emerging pseudo-wire standards even let MPLS networks carry non-IP traffic. MPLS is multiprotocol at last.

Now MPLS is working its way into metro and access networks and even showing up in some private enterprise networks. What's driving this expansion? How will it affect existing services like frame relay and ATM (asynchronous transfer mode)? And what happens when metro-MPLS meets metro-Ethernet?

Three's A Crowd

The spread of MPLS into metro and access networks is driven by familiar business goals: cost reduction and revenue growth. MPLS promises to cut metro-area bandwidth costs by consolidating diverse access traffic. (Can you say "convergence?") It also promises to boost service provider revenue by supporting more robust, scalable Ethernet services.

Today, major service providers support at least two types of wide-area packet services: IP and frame relay/ATM. Access to the two services is provisioned separately. Different lines or TDM channels carry IP and frame relay/ATM packets to their respective backbones. Separate management systems oversee the two services. Now, burgeoning Ethernet services are adding a third packet overlay to metro and access networks. Ethernet's low cost, vast range and high speed make it an attractive alternative to traditional wide-area interfaces for Internet access and IP-VPNs. And Ethernet-based transparent LAN services make it easy for enterprises to stretch their local networks or campus backbones across a wider geography.

Cost-wise, however, three packet overlays—IP, frame relay/ATM and Ethernet—are two too many. Granting each overlay its own TDM channels wastes bandwidth, and managing the overlays separately drives up operating expense. Still, service providers are not about to drop any of the services. As Ray Mota, chief research officer at Synergy Research Group, pointed out, "Frame relay/ATM is a profitable cash cow. Newer services generate valuable revenue, but they're not as profitable yet." The potential to save money by packet-multiplexing all three traffic types into shared channels is one reason that service providers are interested in metro-area MPLS.

Limits To Growth

The other reason is revenue growth. Service providers would like to sell more Ethernet services, both point-to-point and multipoint-to-multipoint. But Ethernet switches, the current basis of most metro-Ethernet services, are Layer 2 bridges, and bridged networks are subject to broadcast storms, Spanning Tree protocol instability and other scaling problems.

Layer 2 virtual LANs present another, more quantifiable challenge. VLANs identify different customers' data streams as they traverse shared lines and switches. But the IEEE 802.1Q VLAN standard accommodates only 4,096 labels, restricting how many customers a provider can serve with one network. The restriction is particularly troublesome if the provider wants to link metro-area Ethernet clouds into a national service. By using LSPs instead of VLANs to segregate customers' Ethernet traffic, MPLS promises to remove the limits to growth.

Multi-protocol At Last

To deliver on these promises—more efficient multiprotocol transport and more scalable Ethernet

Bob Bellman, principal of Brook Trail Research, provides independent product marketing and writing services to the telecommunications industry. He can be reached via email at bob@brooktrail.com.

FIGURE 1 VPLS Provisions A Full Mesh Of LSPs Between PE Devices

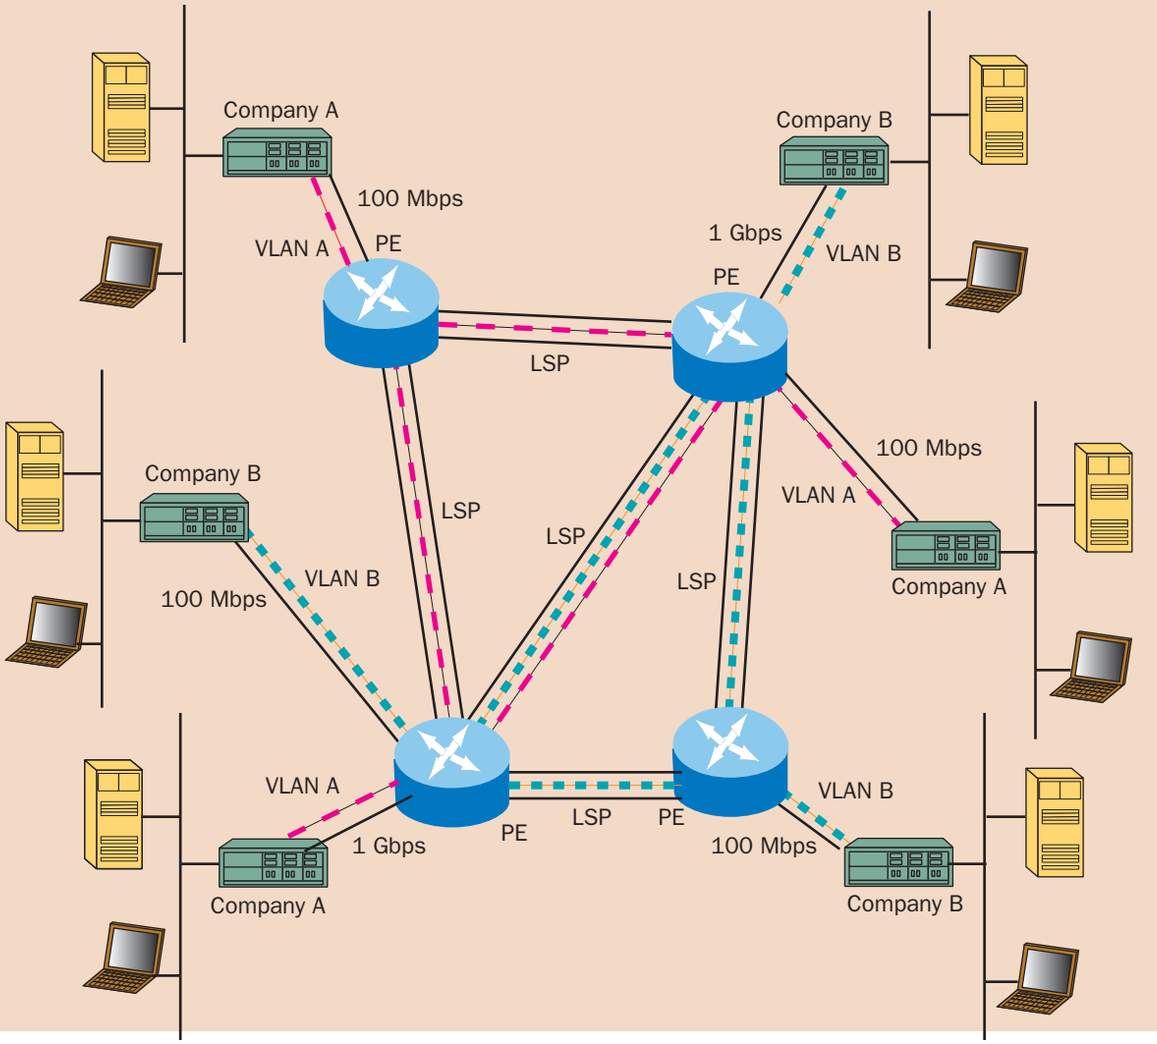

support—the IETF has developed a suite of Pseudo-wire Emulation Edge-to-Edge (PWE3) and Virtual Private LAN Service (VPLS) enhancements to MPLS. Still officially considered drafts, both sets of specifications are supported in prestandard form in many installed devices.

PWE3 (pronounced "pee wee three") defines a mechanism for carrying Layer 1 and Layer 2 traffic in MPLS LSPs. Commonly referred to as the Martini drafts, the PWE3 specs define encapsulations for frame relay, ATM, High-level Data Link Control protocol (HDLC), PPP (Point to Point Protocol), Ethernet and more. PWE3 works at the edges of an MPLS network, adapting Layer 2 frames for transport on input and converting them back to their original form on output. Thus a PWE3 tunnel behaves, at least metaphorically, like an unshared wire.

Using PWE3, a service provider can transport frame relay, ATM and Ethernet frames side-by-side with IP packets over an Ethernet link or a shared TDM channel. The pseudo-wires can carry point-to-point Ethernet across the metro area or provide transparent access to wide-area frame relay/ATM service. Merging diverse traffic types under one statistical multiplexing regime conserves bandwidth, reduces port requirements and simplifies network management.

Metro-MPLS Meets Metro-Ethernet

By applying MPLS VPN techniques to Ethernet services, VPLS eliminates the shortcomings of pure Layer 2 LAN bridging. IP/MPLS routing protocols cope with network faults more reliably than spanning tree protocols, and MPLS labeling lifts the VLAN numbering constraints of IEEE 802.1Q.

For each Ethernet virtual network, VPLS provisions a full mesh of LSPs between the provider edge (PE) devices that serve the customer's locations (Figure 1). Using standard learning bridge techniques, the PE devices switch Ethernet frames between locations based on their MAC addresses. The frames travel in pseudo-wires within the

The big North American carriers aren't rushing to put MPLS in their metro networks

LSPs. Because LSP labels have only local significance, an MPLS network can support an unlimited number of Ethernet VPNs. An enterprise can set up multiple Ethernet virtual networks with different topologies and different QOS parameters: a voice over IP (VOIP) network that reaches all locations, a file-transfer network between manufacturing and warehouses, etc. To simplify wide-area transport, LSPs with shared endpoints and similar QOS parameters can be stacked inside "trunk" LSPs.

With the arrival of PWE3 and VPLS, service providers have several reasons to consider deploying MPLS in their metro and access networks: robust, scalable support for metro and wide-area Ethernet, preservation of profitable legacy services, simplified operation, and cost reduction through bandwidth conservation. Jeremy Chappell, distinguished member of executive staff at Mangrove Systems, estimates "a 50 percent improvement in bandwidth efficiency, especially with Ethernet in the mix, because Ethernet is more bursty than other traffic."

The timing is good too. "ATM and frame relay are on 10-year old platforms," said Chappell. "Service providers are ready to phase out the gear, but they need to preserve the services. Ideally they'd like to replace multiple frame relay/ATM switches with a single edge router." MPLS gives them that option.

Chip Redden, vice president of marketing and product management for Overture Networks, thinks the opportunity for network simplification makes metro-MPLS inevitable. "Eventually, MPLS will replace Ethernet VLANs because the core is already MPLS," said Redden. "It simplifies QOS if everything is MPLS. The adoption of MPLS in metro networks will be driven by SLAs."

New Standards Beget New Products

Not surprisingly, the development of PWE3 and VPLS has sparked several new and enhanced products. Companies like Mangrove Systems and Overture Networks offer devices that sit at the edge of the metro network and consolidate Layer 1 and Layer 2 subscriber traffic into pseudo-wires. If required, similar devices terminate the pseudo-wires at the other end of the metro network and steer subscriber traffic to the appropriate core networks. To compete more successfully with Ethernet switches, companies like Riverstone Networks have added MPLS support to their IP edge routers.

A new breed of MPLS-enabled routers is also emerging. Dubbed multiservice edge (MSE) routers by Synergy Research Group, these devices sit at the edge of a service provider's IP/MPLS network and represent the consolidation of IP edge routers, frame relay/ATM switches and Ethernet switches onto a single platform. Although MSE routers are not available from many vendors yet, Ray Mota of Synergy Research Group expects several IP edge routers to morph into MSE routers, promoting the convergence of both backbone and metro/access networks.

Taking The Plunge

Although committed to MPLS in their core networks, major North American carriers are not leading the charge to metro-area MPLS. AT&T, for example, has had an MPLS wide-area backbone for years, supporting its IP-enabled Frame Relay and other services. AT&T also offers metro-area Ethernet services. But the company is in no hurry to extend MPLS into its metro networks. "We have an evolution plan based on MPLS, including VPLS and PWE3," explained Richard Klapman, group manager for converged packet access, "but none of this is standard yet. We're waiting for standards."

Verizon has a national MPLS backbone that supports pseudo-wires and transparent LAN service. Aggregation routers at the network edge connect to customers via private lines, frame relay, ATM and Ethernet. "MPLS lets everything converge on a single backbone architecture," said Tom Roche, executive director for advanced products and services.

Roche continued: "By connecting all Verizon's exchanges, the MPLS backbone opens opportunities for Verizon that were formerly only for interexchange carriers (IXCs)." Like AT&T, however, Verizon is not rushing to install MPLS outside the core. "We've explored MPLS in the metro," said Roche, "and Ethernet has emerged as the strongest metro technology."

Leave it to more exotic organizations to take the plunge. The Utah Telecommunications Open Infrastructure Agency (UTOPIA) is building an MPLS network based on routers from Riverstone. When completed, the network will bring Fast or Gigabit Ethernet connectivity to approximately 140,000 residential and business subscribers in 14 Utah cities. Service providers will use the wholesale UTOPIA network to deliver voice, data and video services to their subscribers.

Telefónica de España, Spain's leading service provider, is rolling out its Imagenio multimedia broadband service based on VPLS over IP/MPLS. Imagenio will deliver broadcast TV, video-on-demand and Internet access to residential subscribers in major cities across Spain. MPLS-enabled routers from Riverstone will integrate with DSL access multiplexers (DSLAMs) from Lucent Technologies to aggregate broadband traffic onto a converged backbone. An Ethernet virtual network in each metro area will connect subscribers to local video sources.

Metro Sí, Last Mile No

Since both span multiple cities, the UTOPIA and Telefónica examples may not qualify as pure metro-area MPLS. But they do bring MPLS all the way to the service provider edge, stopping just short of the last-mile access circuits. That will

FIGURE 2 Layer-2 Traffic Is Aggregated At The CO And Backhauled To An MPLS Router

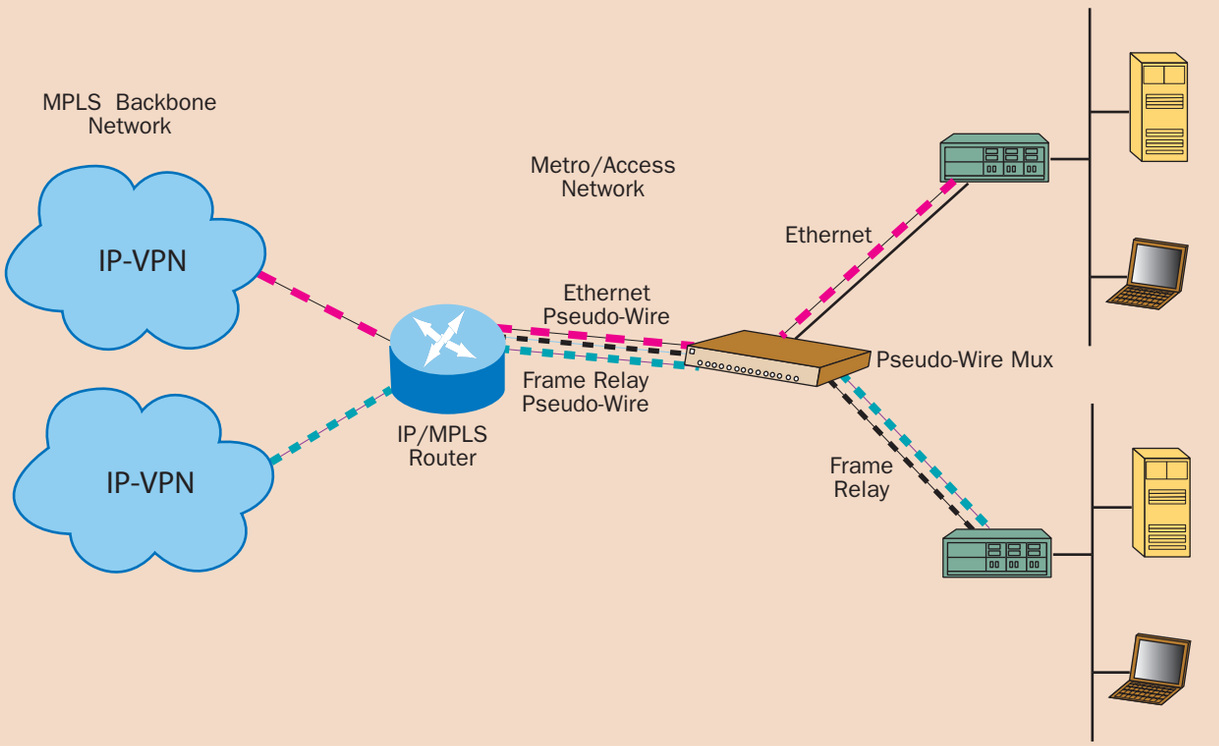

likely remain the outer limit for MPLS for quite a while.

Hardly anyone is talking seriously about extending MPLS to the customer premises. Service providers are sticking with familiar interfaces to the customer: frame relay, ATM, IP/PPP and increasingly, Ethernet. Although the different protocols may share a SONET or DSx access line, they will continue to run over separate TDM channels. Consolidation onto a shared MPLS infrastructure will occur in the central office (CO), behind DSLAMs or cable modem termination systems (CMTSs).

According to Jeremy Chappell of Mangrove Systems, management complexity is one of the reasons that MPLS stops at the CO or sooner. "An MPLS core requires BGP expertise [to manage RFC 2547 VPNs]. The [carrier's] network operations crew can support some number of MPLS routers at the network edge, but they can't support thousands of MPLS endpoints at customer premises." The solution, according to Chappell, is to run Layer 2 rather than Layer 3 connections to the customers and then aggregate the Layer 2 traffic and backhaul it over pseudo-wires to a manageable number of MPLS routers (Figure 2).

Customer perceptions are another reason that MPLS stops at the CO. "Most people see MPLS as a tool, not a service," said Andy Malis, president and chairman of the MPLS and Frame Relay Alliance (MFA) and chief technologist for Tellabs. "They see MPLS in the metro being used to carry IP or Ethernet or frame relay and ATM, not to deliver MPLS services." Plus, no one's eager to suggest that subscribers buy new MPLS-capable CPE.

Demand Puny For MPLS UNI

The lack of interest in last-mile MPLS is epitomized by the fate of the MPLS User-to-Network Interface (UNI). Approved by the MFA in May 2003, the MPLS UNI Implementation Agreement defines an interface for the connection of CPE to a public MPLS network. But so far, the UNI has received little support.

"We haven't seen any demand for the MPLS UNI from carriers," said Azhar Sayeed, Cisco's product line manager for MPLS, "so we haven't built it." Neither, apparently, has anyone else. Andy Malis cites poor timing as the problem. "The MPLS UNI was caught in a chicken-and-egg thing. It came out during the telecom winter. Providers waited for MPLS CPE while CPE vendors waited for MPLS services." Ethernet's ascendancy is another explanation. As Verizon's Tom Roche said: "Customers are settling on Ethernet services. They're familiar with it. There's a comfort factor."

Meanwhile, the MFA has moved on. According to Gary Leonard, director of solutions marketing at Riverstone and the MFA's vice president of marketing, "The MFA is busy on interworking frame relay, ATM and Ethernet over MPLS."

Public frame relay/ATM networks have long

Only the biggest enterprises are implementing private MPLS

been able to aggregate low-speed frame relay traffic from an enterprise's branch offices and deliver it to company headquarters over a high-speed ATM link. Somewhere inside the networks, ATM-to-frame relay service interworking makes the necessary translations.

The MFA's latest efforts will define similar capabilities, but over MPLS LSPs. It will include not only ATM-to-frame relay but also ATM-to-Ethernet and frame relay-to-Ethernet. Enterprises subscribing to MPLS services will be able to exchange traffic transparently between Ethernet sites and frame relay/ATM locations, and they'll be able to change headquarters ATM ports to Gigabit Ethernet without changing anything at the branches.

Private MPLS

Although MPLS is predominantly a service provider technology, a few enterprise companies have built private MPLS backbones—IP/MPLS routers interconnected by private lines. Azhar Sayeed believes that more than 50 Cisco customers have their own MPLS nets. Cathy Gadecki, manager of solutions marketing at Juniper Networks, believes the number for all vendors worldwide is less than 40.

Everyone agrees that only the biggest enterprise networks are implementing private MPLS. "They often have an ATM backbone and just want a 'better' network," said Gadecki: "Better resiliency, more efficiency than ATM, or QOS for voice. They're also looking for 'transparency' so that each department or subsidiary can run its own virtual network over a shared backbone."

Andy Malis concurred. "Enterprises are turning on MPLS for the same reason as telcos," said Malis. "To make optimal use of leased lines between routers and to use MPLS VPNs to partition their backbones for separate departments." Azhar Sayeed cited several examples: A bank using Ethernet to branches via PWE3 over private lines; an airport providing separate VPNs for cargo handlers, airlines, departure control, the FAA, etc.; a financial company using MPLS VPNs to carry traffic from new acquisitions until the businesses are integrated.

At some point, these companies may want a direct interface between their private nets and public MPLS services, perhaps to extend their backbones to locations that don't justify private lines. And at least one VOIP equipment vendor wants its softswitch to be able to signal for LSPs with specific QOS parameters. Such demands could reinvigorate the MPLS UNI.

Alternatively, work is under way to standardize the interface between carriers' MPLS backbones. Once this is accomplished, service providers could then use the same interface for direct MPLS connections to enterprise customers who desire it.

Conclusion

Together, the cost savings of PWE3, the scalability benefits of VPLS, the QOS and reliability features of MPLS LSPs, and the pull of MPLS core networks appear to form an irresistible force. While MPLS won't push out all competing technologies—nothing ever does—it's clearly destined for a major role in future metro and access networks□

Companies Mentioned In This Article

AT&T (www.att.com)

Cisco (www.cisco.com)

Juniper (www.juniper.net)

Lucent (www.lucent.com)

Mangrove Systems (www.mangrovesystems.com)

Overture Networks (www.overturenetworks.com)

Riverstone Networks (www.riverstonenetworks.com)

Telefónica de España (www.telefonica.com)

Tellabs (www.tellabs.com)

Utah Telecommunications Open Infrastructure Agency (www.utopianet.org)

Verizon (www.verizon.com)