MPLS' Newest Application: Layer-2 VPNs

Mary Petrosky

MPLS is neither simple nor a sure thing, but there's lots of heavy betting that it'll succeed. Here's a status report.

At its inception in the mid-1990s, Multiprotocol Label Switching (MPLS) emerged as a fast-packet forwarding scheme to improve router performance. But when the advent of hardware-based routing eclipsed this use, the industry's focus shifted to applying MPLS's connection-oriented characteristics to traffic engineering in large networks. More recently, MPLS has been used for creating Layer-3 (IP-based) virtual private networks (VPNs), based on techniques defined by the IETF in RFC 2547, RFC 2764 and other documents.

Now, another application for the technology has emerged—Layer 2 VPNs. Spurred by interest among service providers, IETF specifications are being defined that spell out how Layer 2 traffic, such as Ethernet, frame relay and ATM traffic, can be transported across an MPLS network. This will allow service providers to accommodate legacy customer traffic—especially lucrative frame relay services—while moving to next-generation IP-oriented network architectures.

The technology also may appeal to IT managers in large enterprises that are used to running their own metro and wide-area networks. For small- and medium-size businesses, network managers could see more service options, including emulated LAN services, as well as higher-speed connections for existing Layer 2 services.

Although no standards have yet been nailed down, several equipment makers have already announced support for the IETF Martini drafts, with more vendor implementations in the works. Bearing the name of Level 3's senior architect, Luca Martini, the Martini drafts define encapsulation and label-distribution mechanisms for transporting frame relay, ATM, Ethernet, High-level Data Link Control (HDLC) and Point-to-Point Protocol (PPP) traffic across an MPLS network. Layer 2 VPN-related documents coming out of the Provider Provisioned Virtual Private Networks (PPVPN) and Pseudo Wire Emulation Edge to Edge (PWE3) working groups support tunneling using IP and L2TP as well as MPLS, but most of the buzz is around MPLS.

Why Layer 2 VPNs?

A VPN is simply a way to provide private communications over a public network infrastructure. For many years, enterprises have purchased data-link layer (Layer 2) connectivity from service providers and created their own Layer 3 infrastructure over those links. In a Layer 2 VPN (L2 VPN), the provider's equipment forwards customer data based on information in the Layer 2 headers, such as a frame relay data link connection identifier (DLCI), an Ethernet MAC address or 802.1q virtual LAN (VLAN) tag.

L2 VPNs are multiprotocol in nature, so they can support both IP and non-IP traffic. They also eliminate the need for service providers to participate in a customer's Layer 3 routing, which can benefit both customers and service providers. For example, customers can hide their routing—

FIGURE 1 Label Stacking In An MPLS Network

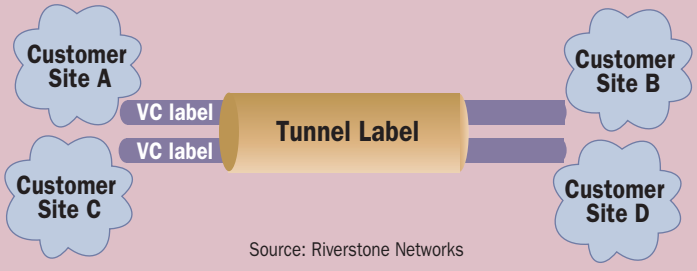

Source: Riverstone Networks

Mary Petrosky is an independent technology analyst based in San Mateo, CA. She can be reached at mary@mpetrosky.com.

FIGURE 2 Frame Relay Tunneled Across An MPLS Network

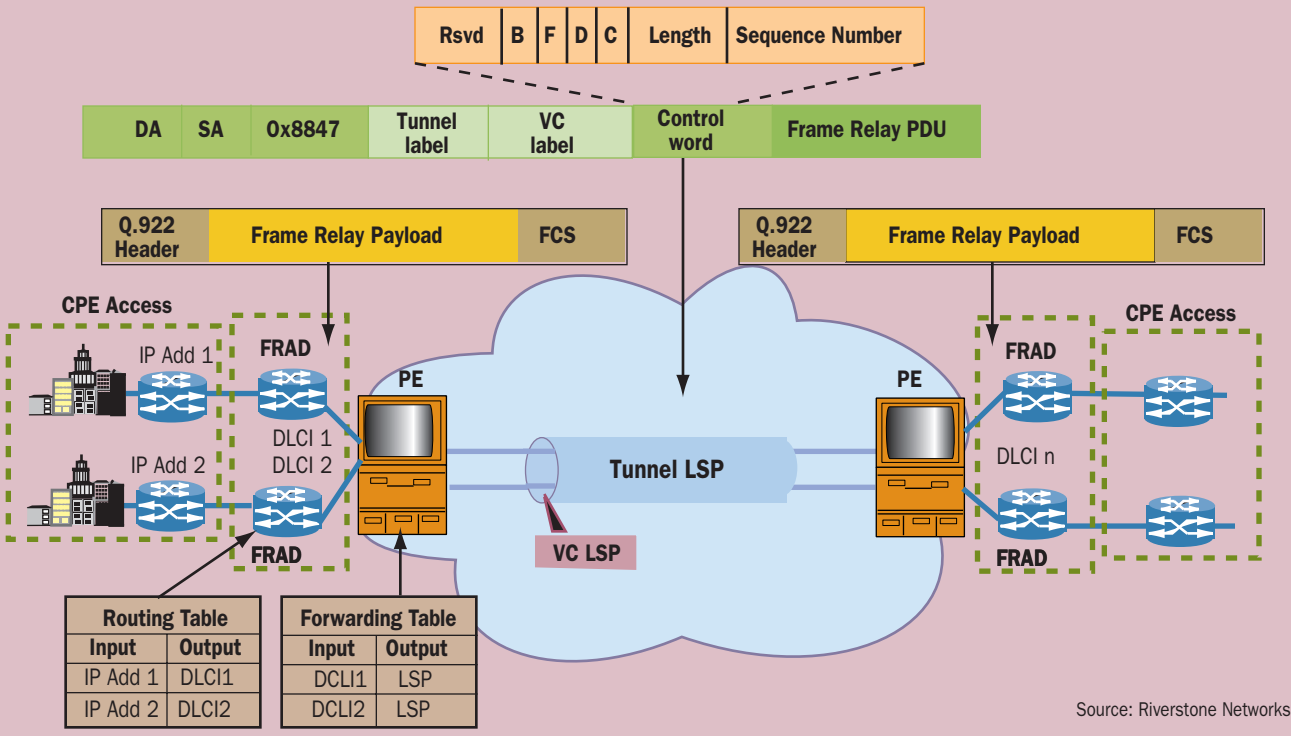

Source: Riverstone Networks

including private IP addressing schemes—from a provider, while providers avoid the complexity of interfacing with a customer's routing infrastructure, putting less strain on edge routers, because there's no need for per-VPN routing tables.

Today, frame relay customers are the bulk of L2 VPN users, with Ethernet-based L2 VPNs (often in the form of transparent LAN services) increasingly popular as more service providers offer Ethernet-based services. Frame relay services are firmly entrenched and demand continues to grow, providing a significant source of revenue that carriers don't want to jeopardize.

However, for many service providers, maintaining—let alone expanding—their frame relay infrastructure is not an attractive proposition. Scaling frame relay access networks, which typically rely on ATM for aggregation, is proving problematic, for example. As vendors shift their product development away from ATM toward IP-based boxes, service providers have found it difficult to buy ATM interfaces above OC-48 speeds, although that will change when the next generation of ATM gear hits the market.

Some service providers also find it onerous to maintain separate network infrastructures for frame relay, ATM, TDM and IP traffic. For example, Cable & Wireless is looking to consolidate its services on an MPLS-based IP infrastructure, and sees MPLS L2 VPN technology as key to that consolidation. Senior director for strategic network planning, David Garbin, notes that "the savings are very significant" in moving from multiple discrete infrastructures to a common infrastructure that can accommodate a range of customer traffic types and services.

Variation On A Theme

Clearly, MPLS is only one option for deploying L2 VPNs. An L2 VPN service can be created over an IP or MPLS backbone; for example, by provisioning point-to-point "virtual circuits" that run through IP or MPLS tunnels.

However, interest in using MPLS tunneling for L2 VPN services is high because providers can leverage MPLS' traffic engineering and fast rerouting; for example, to improve control over network traffic. "MPLS lets you do things that conventional routing protocols, like BGP, don't do, such as routing around congestion," notes Kevin Dillon, director of product marketing at Juniper Networks.

Providers offering Ethernet-based metro area services see MPLS bringing scalability and reliability to Ethernet, better enabling them to support service level agreements (SLAs). In particular, MPLS's connection-oriented capabilities allow for bandwidth guarantees, fast recovery times, customized rerouting and other features traditionally associated with ATM and frame relay.

MPLS also addresses a major shortcoming of Ethernet in the metro and wide area: the limit on the number of VLANs allowed per switch. Although some vendors have developed proprietary work-arounds for this limitation, MPLS L2 VPN techniques offer a standards-based solution

that leverages MPLS's hierarchical label capability (also known as "label stacking"). For example, service providers can associate MPLS labels with customer VLAN traffic and map VLANs to MPLS label switched paths (LSPs). (Note: Martini encapsulation accommodates both VLAN-tagged and untagged Ethernet frames.)

Economics are another reason MPLS-based L2 VPN approaches are considered to be attractive relative to traditional frame relay and ATM-based L2 VPNs. The cost of packet-based equipment is extremely competitive, which explains why many believe that Ethernet will blossom in the metro area. The wide availability of MPLS-based L3 switches and routers—from enterprise class equipment to core routers—has also led to competitive pricing. But perhaps the biggest competitive driver for using MPLS is the fact that it brings multiservice capabilities to IP networks, offering the possibility of consolidating services onto a single infrastructure.

MPLS L2 VPN Mechanics

IETF efforts to define mechanisms for creating MPLS-based L2 VPNs are predicated on two key concepts. First, with MPLS, it's possible to create a tunnel as an LSP using label switching rather than network-layer encapsulation. Second, it's possible to use control protocols, such as the MPLS Label Distribution Protocol (LDP) and BGP, to set up emulated virtual circuits (VCs) that carry the Protocol Data Units (PDUs) of Layer 2 protocols across a network. A number of such emulated VCs can be carried in a single tunnel so long as the PDUs are encapsulated.

The Martini approach uses MPLS label stacking to support separate VC labels and tunnel labels. The VC label identifies the VPN, VLAN or connection at the end point, while the tunnel label determines the path a packet takes through the network. Label switch routers (LSRs) in a network core use the tunnel label for forwarding, and the egress LSR uses the VC label to determine how to process the frame. Figure 1 (p. 29) provides a high-level illustration of this use of label stacking.

There are two Martini drafts, although they're often simply referred to as "Draft Martini." One, the encapsulation draft, specifies the emulated VC encapsulation for frame relay, ATM (AAL5 and cell), Ethernet (native and VLAN tagged), HDLC and PPP. It also defines a demultiplexer field, which is used in an MPLS environment to distinguish individual emulated VCs within a single tunnel.

In addition, this draft specifies a Control Word, whose functions include preserving the sequence of frames, padding small packets to meet minimum packet sizes and carrying control bits from the original L2 frame header. Since it isn't always necessary to transport the L2 encapsulation across the network, an ingress router may strip the original L2 header from a frame. The egress router would use the Control Word to reproduce a frame's L2 header.

The second Martini draft defines the label distribution procedures needed to transport encapsulated L2 PDUs across an MPLS network. Currently, Martini specifies LDP for setting up the tunnel LSPs. IETF members also are exploring the use of BGP to distribute label blocks and mappings to frame relay DLCIs, ATM virtual channel identifiers (VCIs) and other identifiers. Defined in the Kompella draft "draft-kompella-mpls-l2vpn-02.txt," this approach has its share of supporters, including Juniper Networks. Some vendors, including Tenor Networks, plan to support both LDP and BGP mechanisms for L2 VPN set up.

Figure 2 illustrates the way a Martini-based L2 VPN would work. In this case, frame relay traffic from a customer site is encapsulated for transport across an MPLS metro network. IP traffic from a customer's site is passed through a frame relay access device to a router at the service provider's network edge. There the traffic is encapsulated per the Martini draft and assigned to a virtual circuit (VC) for transport across a tunnel LSP. At the egress router, the encapsulation is stripped off, the frame relay header and frame check sequence (FCS) is recreated and the frame is forwarded to the destination customer site.

Although the encapsulations vary for each L2 traffic type, the basic process illustrated in Figure 2 is the same for ATM, Ethernet, HDLC and PPP.

Commitment Despite Flux

As noted above, the PPVPN and PWE3 are the key IETF working groups defining standards for L2 VPNs. Roughly speaking, the PPVPN group is focused on specifying multipoint VPN technologies that service providers can provision for enterprise customers. In contrast, the PWE3 working group is focused on defining technologies for point-to-point links that extend from edge to edge in a provider's network, so technologies that come out of this working group are likely to be used by service providers internally on their networks.

The Martini drafts pre-date the PWE3 working group, which was only formed in March, 2001. Level 3's Luca Martini, who co-chairs PWE3, expects key features of his drafts to be published in the near future as Experimental RFCs, to reflect their actual implementation in the market.

Indeed, despite the flurry of activity in both the PPVPN and PWE3 working groups, nearly a dozen vendors have committed to or already implemented the Martini drafts, and several service providers expect to deploy it in their networks beginning next year. "Martini is embedded in so many people's consciousness, it's all but a standard," said Tim Wu, technical marketing director for Riverstone Networks.

Steve Vogelsang, co-founder and vice president at Laurel Networks, concurs. He expects that much of the content of the Martini drafts will be

Nearly a dozen vendors have, or plan to implement the Martini drafts

Carriers have also begun to test software based on the Martini drafts

TABLE 1 Rollout Plans For L2 VPN Gear

Vendor	Drafts Supported	L2 Technology Supported	Availability
Alcatel	Fischer encap and transport (PNNI)	ATM AAL5 and cell	Q4 2001
	Martini encap and transport (LDP)	Ethernet	Q2 2002
Atrica	Martini encap and transport (LDP)	Ethernet	Q1 2002
Cisco	Martini encap and transport (LDP)	Ethernet	Nov. 2000
		ATM AAL5	Aug. 2001
		Frame Relay,	Q1 2002
		ATM cell, PPP, HDLC	Q2 2002
Extreme Networks	Martini encap and transport (LDP)	Ethernet	June 2001
Foundry Networks	Martini encap and transport (LDP)	Ethernet	Nov. 2001
Laurel Networks	Martini encap and transport (LDP)	Ethernet, ATM AAL5 and cell, frame relay	Q4 2001
Nortel Networks	Koleyni/Fischer encap and transport (PNNI)	ATM AAL5 and cell	Q3 2001
	Martini encap and transport (LDP)	Ethernet	Q1 2002
Riverstone Networks	Martini encap and transport (LDP); Lasserre	Ethernet	June 2001
		ATM cell, frame relay	Q1 2002
Tenor Networks	Martini encap and transport (LDP); Kompella (BGP)	Ethernet, frame relay	Q2 2002

rolled into the PWE3 as the basis for that group's work going forward. Rather than major changes to Martini, Vogelsang expects PWE3 to provide a forum for broader analysis of the drafts and for ironing out issues that arise relating to specific L2 technologies.

Multiple drafts focused on ATM and Ethernet have already popped up. For example, the so-called Koleyni, Fischer (a revised version of Koleyni) and Brayley drafts outline ATM-specific requirements and encapsulation formats and semantics. The Kamapabhava draft, which emerged this fall in the PWE3 working group, focuses on frame relay. In many cases, the Martini drafts are serving as the starting point for further extensions. For Ethernet, for example, the Lasserre draft proposes extensions to Martini to support multipoint-to-multipoint connectivity, allowing for broadcasting and multicasting across VPN and TLS connections.

Products And Services In The Pipeline

Cisco claims to be the earliest supporter of the Martini drafts, having developed an ATM AAL5 encapsulation implementation for field trials 12 months ago. "Cisco's very much behind Draft Martini," noted Azhar Sayeed, IP MPLS Manager in the company's IOS Technologies Division, adding that Cisco worked closely with Level 3's Martini on this early implementation.

Among other early supporters of the Martini drafts are Extreme Networks and Riverstone Networks, both of which announced support for these technologies in mid-2001. At N+I in September, Cisco, Extreme, Foundry, Laurel and Riverstone participated in interoperability testing of Ethernet over MPLS across an MPLS network compliant with Martini encapsulation. Vendors forwarded L2 traffic using LDP over LDP as well as LDP over RSVP.

This past fall, Laurel Networks and Foundry announced their support for L2 VPN services based on Draft Martini. Atrica, Juniper, Nortel Networks, TiMetra and Tenor Networks are also closely tracking Martini and related drafts, with plans to implement these technologies in the coming year. Table 1 summarizes some of these vendor implementations and plans. (Note: The table encompasses MPLS L2 VPN-related technologies only; in many cases, vendors are supporting additional MPLS protocols, such as RSVP-TE and CR-LDP.)

A handful of service providers have already begun evaluating and testing Martini-based code. Likewise, several large enterprises, particularly in the financial services and government areas, have shown interest in the technology. Having spearheaded development of the Martini drafts, it's no surprise that Level 3 is leading the deployment charge. "We are very much in the implementation process and have plans for rolling out services in January," said Luca Martini.

Cable & Wireless is also quite interested in MPLS L2 VPNs and the Martini approach, according to Chris Liljenstolpe, senior director of network technology. Cable & Wireless began bringing equipment into its labs last year and expects to make final equipment selections by mid-2002 and to begin field trials in the second half of this year. "I'm impressed with how quickly vendors have gotten interoperable implementations out the door," Liljenstolpe said, adding that

the service provider already has several large customers interested in L2 VPN services.

Riverstone's Wu indicated that Hong Kong's incumbent carrier, Pacific Century Cyberworks, and British service provider Storm Telecommunications are among the service providers that have already deployed Riverstone's Martini L2 VPN technology. He also noted that several metro service providers, including IntelliSpace and Telseon, are interested in it as well.

Indeed, co-founder and chief technical officer, Carlo Lalomia, noted that IntelliSpace's goal is to deploy MPLS throughout its network, even down to the customer premise. "Today we have a fully routed architecture. In the future, we're looking to use Martini-based L2 VPNs as a method to quickly provision new customers and provide them with additional services." IntelliSpace is currently testing MPLS implementations from several vendors with an eye toward delivering services in mid-2002.

Telseon is most interested in the Martini Ethernet encapsulation, according to Jay Gill, director of product management for the company's Gigabit Service. Having standards-based encapsulation will allow Telseon to transport customer traffic transparently, without interfering with the internal operation of the company's network. The Martini technology will allow Telseon to offer customers an Ethernet-based transparent LAN service with support for VLAN trunking, and give customers the flexibility to use either Ethernet switches or routers at their premise. Telseon plans to roll out services based on the Martini technology in Q1 of this year.

Yipes Communications is also keen on the Martini approach. "We're interested in any technology that allows for L2 services in a scalable, cost effective, efficient fashion," said CTO Kamran Sistanizadeh. Currently, Sistanizadeh is watching the Martini and Kompella drafts, with a particular focus on the Martini encapsulation and transport drafts as a means for Yipes to hand off its metro and regional traffic to upstream providers. In addition, Sistanizadeh sees benefits for its customers: "With MPLS L2 VPNs you can create native L2 solutions for customers."

The Bottom Line

MPLS-based L2 VPNs are attractive to a wide range of service providers because they promise to help providers make money as well as save money. Transport-oriented carriers, such as RBOCs and ILECs, that are used to circuit technologies and have limited experience running fully routed networks, find the MPLS L2 VPN model a comfortable fit. On the opposite end of the spectrum, providers focused on Internet connectivity and IP-based services can use the Martini technology to trunk customer traffic across their networks without worrying about the type of traffic being transported.

While the emphasis in recent years has been on IP and IP-related services, the fact remains that many service providers derive a significant portion of their revenue from data and voice services based on legacy technologies such as frame relay, ATM, and time division multiplexing (TDM). Greenfield and established service providers alike want to make money from these traditional services even as they build out their IP infrastructures. Service providers such as Level 3 and Cable & Wireless see MPLS-based L2 VPNs as key to their ability to consolidate their network architectures around IP and MPLS while supporting existing customer frame relay and ATM traffic and potentially expanding into Ethernet transport.

Clearly, development of MPLS L2 VPN standards is in the early stages. Early adopters of Martini-based solutions face the classic problems all early adopters face, including a lack of solid provisioning and management tools. However, momentum is building for MPLS L2 VPNs built around the Martini approach, because the solution addresses real service provider and enterprise customer requirements. Regardless of how widely it's ultimately deployed, the Martini approach has already found a market. □

Carriers want to keep making money from their traditional services while building out new IP infrastructures

Companies Mentioned In This Article

Alcatel (www.usa.alcatel.com)

Atrica (www.atrica.com)

Cable & Wireless (www.cwusa.com)

Cisco (www.cisco.com)

Extreme Networks (www.extremenetworks.com)

Foundry Networks (www.foundrynet.com)

IntelliSpace (www.intellispace.net)

Laurel Networks (www.laurelnetworks.com)

Level 3 (www.level3.com)

Nortel Networks (www.nortelnetworks.com)

Juniper Networks (www.juniper.net)

Pacific Century Cyberworks (www.cyberworks.com.sg/main.htm)

Riverstone Networks (www.riverstonenet.com)

Storm Telecommunications (www.stormtel.com)

Telseon (www.telseon.com)

Tenor Networks (www.tenornetworks.com)

TiMetra (www.timetra.com)

Yipes Communications (www.yipes.com)