



## MPLS VPNs:

# THE REAL DEAL

Multiprotocol Label Switching is ready to revolutionize service-provider routing. But can this developing technology affect your company? You bet. BY BRUCE BOARDMAN

In our July 10, 2000, primer on MPLS, we explained that Multiprotocol Label Switching gets packets to their destinations efficiently by creating paths through a network, similar to ATM and frame relay (can you say *QoS?*), while retaining the flexibility of IP (see “MPLS: A New Traffic Cop for Your WAN,” at [www.nwc.com/1113/1113ws2.html](http://www.nwc.com/1113/1113ws2.html)). Recently, you may have heard a buzz about MPLS VPNs. For many enterprise customers, having an MPLS circuit referred to as a VPN is confusing. After all, an MPLS VPN isn’t encrypted—it doesn’t even go across the Internet. What’s more, the MPLS standards have been in the oven for years but are still not quite baked, and the few products around are for service providers. It’s hard for most enterprises to care about what seems to be an immature, insecure technology.

Although it seems misleading, an MPLS circuit is accurately referred to as a VPN. While an IPsec VPN creates a circuit across the Internet, MPLS creates circuits across service providers’ networks. Are MPLS circuits less private because they aren’t encrypted? Yes and no. Can enterprises

ignore MPLS VPNs because they are a service-provider play? No. Most enterprises aren’t going to implement MPLS, but increasingly service providers will provision enterprise WAN circuits using MPLS. Managing an MPLS service will likely require an end-to-end SLA (service-level agreement), just as ATM and frame relay services need today. Future cool is the possibility of self-provisioning WAN circuits on demand using MPLS VPNs.

## WHAT’S YOUR ADDRESS?

All MPLS applications stem from the new address space created by the MPLS label. The MPLS label, while co-existing with Layer 2 and Layer 3 infrastructures, makes traffic manipulation deterministic. This is similar to the tagging in Ethernet VLANs, but it applies to multiple protocols. MPLS is very important to service providers because of its speed and traffic-engineering capabilities; those providers are deploying it to leverage IP in the core of their networks while maintaining control over how bandwidth is used in that core. The combo of IP flexibility with traffic-engineering capabilities is the beauty of MPLS. In addition, service providers continue to use frame relay and ATM for access as they move to newer last-mile access protocols, like MAN Ethernet. This blend of old and new provides value for existing equipment and a way to migrate customers.

The MPLS label is an index to a route, or LSP (Label Switch Path), so rather than running a longest match on the destination IP address, the label indexes at each hop the next hop behavior. In addition to being faster, MPLS offers a couple other advantages: Layer 3 address transparency, privacy, scalability and clearly defined management boundaries (for an example of a Layer 3 MPLS VPN that illustrates this last point, see “Layer 3 MPLS VLN,” page 106).

## MPLS VPN STANDARDS

Confusing the role of MPLS VPNs is the current lack of a well-defined, implemented standards. More than 50 MPLS standards and drafts are floating around; 15 of these are for MPLS VPNs (for the whole shebang, go to [www.mplsrfc.com/standards.shtml](http://www.mplsrfc.com/standards.shtml)).

MPLS VPN standards are a subset of the same group of proposals that make up the MPLS standard. There are, however, MPLS VPN working groups specifically proposing



## Glossary

Martini draft	An Internet draft written by Luca Martini and many others. It defines how MPLS can be used to support Layer 2 VPN services, such as Ethernet, frame relay and ATM.
RFC 1918	Describes address allocation for private networks.
RFC 2547bis	Describes a method by which a service provider may use an IP backbone to provide VPNs for its customers. MPLS is used to forward packets over the backbone, and BGP is used to distribute routes over the backbone.
VPN Routing and Forwarding Table (VRF Table)	The table that an MPLS label references to determine next hop. The VRF table used for lookup is based on the interface on which the packet arrives.

## WORKSHOP

IETF standards. The Provider Provisioned Virtual Private Network (PPVPN) group, for example, is chartered with creating a VPN framework to suggest best practices for MPLS VPN security, scalability and manageability ([www.ietf.org/html.charters/ppvpn-charter.html](http://www.ietf.org/html.charters/ppvpn-charter.html)).

The three basic types of VPNs the PPVPN group is considering are MPLS BGP (Border Gateway Protocol) VPNs, MPLS Virtual Routers and MPLS Layer 2 VPNs. The PPVPN working group is coordinating with the Pseudo Wire Emulation Edge to Edge (PWE3) working group, which is creating standards for tunneling end-to-end connections through ATM and MPLS networking fabrics at Layers 1 and 2 ([www.ietf.org/html.charters/pwe3-charter.html](http://www.ietf.org/html.charters/pwe3-charter.html)).

Finally, Layer 2 VPNs are specified in the Martini draft now in the IETF PWE3 working group. The idea is to tunnel Ethernet, frame relay, ATM and PPP (Point-to-Point Protocol) within MPLS. The PWE3

is working on other similar standards, but Martini is getting the most attention from service providers.

### GOT THAT?

Adding to the confusion, the term *VPN*, when applied to MPLS, is used differently from its common meaning. VPNs have come to be defined as encrypted tunnels that ride over Layer 3 pro-

ocols. The encryption and data of the encrypted packet unreadable and, thus, private. This encrypted payload is placed in another packet that carries it across a network. Upon arrival, the encrypted packet is removed and unencrypted.

MPLS VPNs are also PVCs (private virtual circuits), like IPsec or PPTP (Point-to-Point Tunneling Protocol) VPNs, but that's where the similarity ends. In an MPLS VPN, privacy doesn't come from encapsulation or encryption. In fact, there is no encryption at all. Privacy comes from

segregating packets based on their MPLS labels. Traffic for a particular label is read only by the LSRs (Label Switch Routers) along that LSP. Normal IP routing methods are not applied within the MPLS fabric—only the MPLS labels are read to deliver traffic. If this makes you nervous, consider that this level of security is equal to

## MPLS VPN PRIVACY DOESN'T COME FROM ENCAPSULATION OR ENCRYPTION. IN FACT, THERE IS NO ENCRYPTION AT ALL.

that of existing Layer 2 protocol links in that the data flowing over ATM or frame relay PVCs is also unencrypted. For the truly paranoid, there's no law barring encryption of the packets to which MPLS headers are attached.

### AN MPLS BGP VPN IN ACTION

Traditional WAN circuits as offered by service providers are made up of Layer 1 and Layer 2 protocols. This means that enterprises with multiple offices purchase Layer 1 access connections, like T1 lines, from each office into a

service provider's network; a Layer 2 protocol, like frame relay or ATM, is then used to traverse from the edge routers supporting the T1s through the service provider's core network and onward to the remote offices' edge routers.

An MPLS circuit is the same. At each enterprise location some type of Layer 1 and Layer 2 access circuit

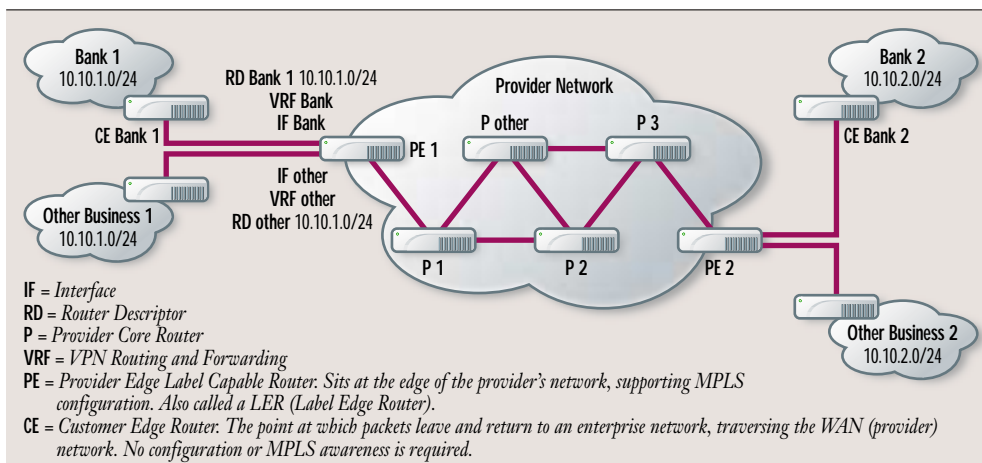
is still needed. But instead of traversing the core of the service provider's network using ATM or frame relay, an MPLS VPN is built using MPLS labels.

Lets look at how an RFC 2547bis ([www.ietf.org/internet-drafts/draft-ietf-ppvpn-rfc2547bis-01.txt](http://www.ietf.org/internet-drafts/draft-ietf-ppvpn-rfc2547bis-01.txt)) MPLS BGP-VPN works (see "Layer 3 MPLS VLN," at left). Note that the types of Layer 1 and 2 connections used are not important for this example.

The players in our little comedy are, from right to left, the enterprise, or CE (customer edge), routers, which terminate Layers 1, 2 and 3 protocols at the edge of the customer's network; the PE (provider edge) router, which is upstream from the CE router at the edge of the provider's core network, doing most of the work; and finally the P (provider) core routers, supporting MPLS LSPs.

The CE router is attached by either a static or an IGP (Interior Gateway Protocol) routing protocol, including RIP, OSPF, ISIS (Intermediate System to Intermediate

## LAYER 3 MPLS VPN



System), EGP (Exterior Gateway Protocol) or BGP. What's important here is that the routing protocol used is *not* important. For the customer, this means no coordinated changes or even access beyond basic Layer 3 connectivity with the PE router. For the service provider, this means no messing with hundreds of customer routers to provision or maintain network circuits. This is the defined management boundary.

On the PE router a VRF (VPN Routing and Forwarding) table is configured for a particular interface or subinterface. This is a routing table and the basic building block of an MPLS BGP VPN. Routes learned from the attached CE router are populated in the VRF table; in our example, CE Bank 1 has advertised to PE-1 10.10.1.0/24, which is inserted in VRF-Bank.

The entry in VRF-Bank includes the original IPv4 route with a prepended RD (route descriptor). The fields in the RD ensure that the route is unique and should include include a public ASN (autonomous system number). This combination of IPv4 and RD is referred to as the VPN-IPv4 address. It is most common to create a VPN-IPv4 address, but IPv6 and IPX are also considered in RFC 2858.

PE-1 distributes the route in the VRF using a version of BGP: MP-BGP (Multiprotocol BGP). MP-BGP supports an extended community attribute field, enlarged to 32 bits, from 16. MP-BGP is backward-compatible, but standard BGP is of course unable to distrib-

ute VPN-IPv4 MPLS BGP VPN addresses.

An attribute referred to as an *Export Target* is set within the VRF-Bank on PE-1. This attribute determines which target PE routers will receive the BGP-distributed VPN-IPv4 route listed in VRF-Bank. Likewise, on PE routers that are to accept the route as a defined destination to CE Bank 1, an import target must be set in their coinciding VRF tables. In our example, that target is PE-2 VRF-Bank. Once the route is distributed and accepted, communications can begin.

Note that the PE routers maintain only the CE routes that are defined in their VRF tables, not the routes for all CE routers in the network. Likewise, the P routers are aware of only

## MPLS WILL LOWER SERVICE PROVIDER COSTS, AND EVENTUALLY THE SAVINGS WILL TRICKLE DOWN TO THE ENTERPRISE.

the PE routes, not the VPN or CE routes. This reduces complexity.

MPLS supports a label stack, which creates a hierarchy of LSPs. When a packet leaves CE Bank 1 destined for CE Bank 2, the VRF route lookup on PE-1 uses VRF-Bank 1, based on the interface on which that packet arrived. When a match is found, an MPLS label that relates to the next-hop PE device is inserted into the stack—in this case, PE-2. PE-1 then places on top of the stack the label that specifies the next hop in terms of the next core P router, and the packet is forwarded to P-1. P-1 examines the label, rec-

ognizes the next hop, pops the top PE-1 label and places a label to the next-hop P-2 router, leaving the PE-2 label on the stack.

This continues until the final P router, which in our example is P-3. P-3 pops the P-2 label, examines the PE-2 label, recognizes that PE-2 is directly connected and forwards the packet to PE-2. PE-2 removes the labels and forwards a native IPv4 packet onto the interface of CE Bank 2. The reverse works the same way, except different labels are used. This creates a different LSP, which may very well travel the same path, though it doesn't have to. You can create a meshed or hub-and-spoke network by manipulating the route import and export targets.

Notice that the packets


to carry overlapping nonunique RFC 1918 addresses without requiring NAT, BGP extensions are required. BGP extensions resolve multiple IPv4 addresses, like 10.10 or 192.168. Because the IP prefixes are not used to route, so as long as IP prefixes are not duplicated within a VPN domain, the MPLS VPN cloud has no notion or reliance on the routed prefixes.

**SHOW US THE MONEY**  
WANs cost enterprises big bucks, but there is some hope that MPLS will lower service providers' costs, and eventually the savings will trickle down to the enterprise. If lower WAN pricing isn't in store, perhaps we'll see a shorter, easier provisioning process—or even self-

in our example are moved across the provider cloud based on the shimmed label. The IP address isn't referenced, so NAT (Network Address Translation) and private IP routing aren't necessary. This means that networks with duplicate RFC 1918 addresses, like 10.10.0.0 and 192.168.0.0, can be routed across the provider backbone.

**M**PLS BGP VPNs use BGP to distribute LSPs, which are the instructions that define how packets hop through the network and decide on the route based on IP prefixes.

For MPLS BGP VPNs

provisioning to make WAN circuits available just in time, instead of all the time. These benefits will likely not happen for years, however, because service providers have just begun creating MPLS infrastructures. But keep your eyes on MPLS. 

*Bruce Boardman is executive editor of NETWORK COMPUTING, testing and writing about network management and systems. He has 12 years' IT experience managing networks and distributed computing for a financial service provider. Send your comments on this article to Bruce Boardman at bboardman@nwc.com.*