

# What's Your Flavor of Networking?

By Ken Van Orman

*We don't recommend sipping the Martini in this case, but using it or 2547 will help your service provider develop potentially new sources of revenue and reduce network operating overhead.*

What do a cocktail drink, a seemingly nondescript number and networking have in common? Well, nothing really, but Martini and 2547 are commonly used names to describe two emerging VPN solutions.

Martini is named after Luca Martini, an employee of Level 3 Communications, who is the primary author of two IETF drafts describing a method for transporting Layer 2 frames over MPLS. 2547 is short for another IETF draft for RFC 2547bis, which outlines a method for supporting a Layer 3-based VPN solution.

Virtual private network solutions continue to contribute to a large portion of service provider revenue. By offering innovative services, reducing operating expenditures and reducing capital expenditures, service providers can continue to maintain and even expand profitability in this lucrative market area. Martini and 2547 are two different approaches for tackling these issues.

In general, these types of VPN services are known as provider provisioned VPNs (PPVPN), where the service provider is responsible for provisioning and managing the VPN. In contrast, there exists another class of VPNs used by enterprises. These involve tunneling mechanisms that connect enterprise-owned devices over some public infrastructure. IPSec, PPTP and L2TP are typically used in these solutions. We will focus on PPVPNs in this article.

## The Evolution of VPNs

The ability to connect computers to share data is a concept almost as old as the computer itself. As companies expanded regionally, nationally and globally, the need emerged to provide connectivity between remote offices. This connectivity was initially provided via dedicated leased lines over a time division multiplexed infrastructure. Companies purchased these connections between their remote offices and "owned" all of that line's bandwidth all of the time. These remote networking solutions were costly and time consuming to provision.

To improve the cost model of remote networking solutions, a method of efficiently sharing networking infrastructure evolved. These new technologies employed virtual circuits (VCs), rather than dedicated circuits, to better "share" network resources. As a result, the virtual private network (VPN) emerged. First X.25, then Frame Relay and finally ATM appeared as Layer 2 VPN solutions of this type.

A growth period of VPN services based on these Layer 2 technologies emerged, but new problems of cost and scalability appeared. In addition, emerging IP data services were often built on separate networks and were not able to leverage the existing Layer 2-based infrastructure.

## Enter the MPLS-based VPN

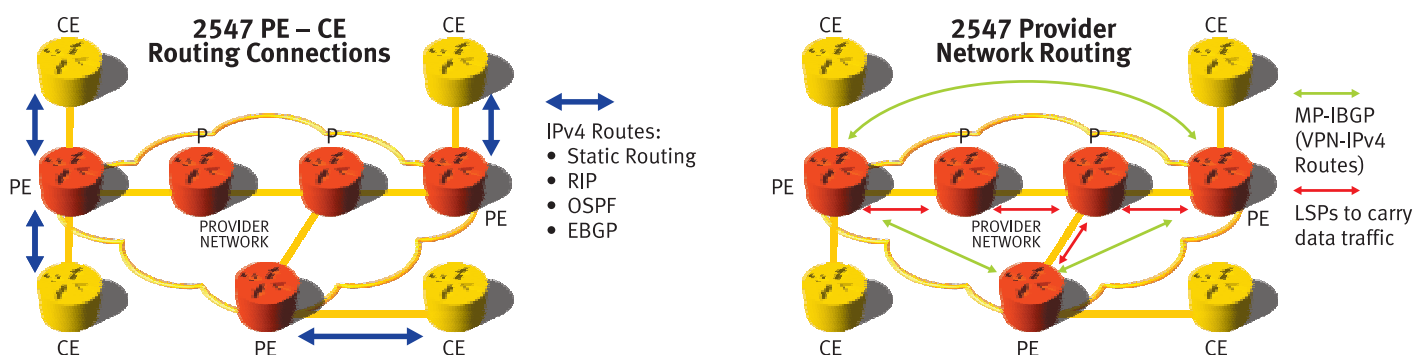
Out of these limitations, yet another type of VPN – the MPLS-based VPN – was created. The two prominent methods are known as Layer 3 and Layer 2 MPLS PPVPNs. More specifically, the Layer 3 PPVPN is based on the RFC 2547bis draft (BGP/MPLS VPN) and the Layer 2 PPVPN is based on the Martini drafts.

The common thread through the two implementations is an MPLS-based core. The MPLS-based core is used as a scalable, tunneling and aggregation mechanism for the VPN traffic. The core routers, also known as provider or P routers, do not connect to any customer devices (also known as customer edge or CE devices). The job of connecting customers to the provider network and maintaining the VPN is left to the provider edge (PE) routers. This allows the P routers to be removed from the internal workings of the VPN and to be used purely as MPLS packet forwarding devices. The direct result of this is a scalable and manageable core network implementation.

The other advantage of an MPLS-based core is the ability to converge IP services networks with VPN networks. MPLS supports forwarding of native IP traffic as well as Layer 2 or Layer 3 VPN traffic. By building one network, service providers can minimize the costs associated with building and supporting separate network infrastructures.

Now let's take a look at the two technologies, separately. BGP/MPLS VPNs offer a Layer 3-based service to the customer. The CE router connects to the PE router through some type of routing mechanism. This can be RIP, OSPF, EBGP or static routes. To provide a means of privacy between VPNs, the PE router "reserves" a portion of its resources for this customer connection to hold the routes for the customer's VPN. To achieve economies of scale, there needs to be many such customer connections on any given physical port, with each having its own "reserved" piece of the router.

The job of propagating these customer routes across the network is left to



# 2547 / MARTINI

Multiprotocol BGP (MP-BGP). MP-BGP peering sessions are established between PE routers and VPN-IPv4 routes are shared. VPN-IPv4 routes are customer routes that have been converted to a special route that carries additional information. The additional information is used to identify which VPN the routes belong to and how to get traffic to those destinations.

That's pretty complex! Now how does that all translate into forwarding data traffic? Lets take a look. PE routers forward traffic by matching incoming packets with entries in that VPN's routing table. Traffic coming and going to connected CE routers is native IPv4, while traffic traversing the core makes use of MPLS for forwarding. Since P routers only know how to forward MPLS labeled packets and know nothing about the VPN traffic, the goal is to get the VPN traffic across the core to the egress PE router using MPLS. The egress PE router is directly connected to the destination CE router and will know how to get the traffic to that CE.

## The Martini Implementation

This is how a Martini implementation works: Unlike a 2547 network, a Martini network is designed to deliver a Layer 2 service to the customer; no routing between CE and PE is required. In fact, a Martini network doesn't care what the customer uses as a Layer 3 protocol. Some Layer 2 services Martini is designed to handle are ATM, Frame Relay and Ethernet.

So if IP routing isn't used, how do we get traffic from customer site to customer site? By setting up a new kind of virtual circuit – a Martini virtual circuit. The connection offered to the customer is still a traditional

Layer 2 circuit. However, this circuit is translated into a Martini virtual circuit (VC) by the ingress PE router. These Martini VCs are provisioned in the ingress and egress PE routers for that VC and signaled to each other via a modified version of LDP. This version of LDP includes specific information for describing the characteristics of the Martini VC. The many Martini VCs provisioned between PEs are tunneled and aggregated via traditional MPLS tunnels of the core network.

Therefore, anyone familiar with traditional IP networking can see some challenges in developing and testing a PE router implementation of either technology. Firstly, we can see there are potential issues with control plane load and activity that conventional routers are not designed to handle – most traditional routers aren't designed to support hundreds of routing sessions, much less thousands that a 2547 implementation is expected to handle. Secondly, there are some new protocols, new types and ways to handle routes and new ways to forward traffic.

When working with new protocol implementations, one of the first steps is to perform conformance and functional testing. For 2547, one will need to verify that customer routes are handled properly and propagated to other PE routers as VPN routes. In turn, incoming VPN routes must be filtered according to policies and propagated as customer routes. Can the implementation do this for RIP, OSPF and EBGP?

For Martini, VCs must be provisioned and signaled with the proper parameters for all supported Layer 2 services. What about the functional testing of data packet forwarding for both protocols? Are the data

packets properly encapsulated and de-encapsulated when they are moved across the core?

## Testing Challenges for Providers, NEMs

Next, capacity issues must be addressed for the control plane and the data plane:

- How many customer connections can a PE handle? How many routes per customer can be handled?
- What about MP-BGP peering sessions and MPLS tunnels on the core side? How many of these are supported?
- Are there different numbers of peers supported for the different customer routing protocols?
- How many Martini LDP peering sessions can a PE support? How many VCs can each PE handle?
- For data traffic, how quickly and accurately can data traffic be forwarded?
- Does VPN or VC leakage occur?
- What happens to data traffic when new VPN routes or VCs are installed or removed? What about when a new VPN site is added?
- What about a whole new VPN with multiple sites?
- How long does it take for new sites or VPNs to be provisioned?

Answers to these capacity questions will help a service provider determine how economical and efficient a network built with the device under test will be. In addition, network equipment manufacturers that have answered these questions will be better positioned when service providers ask how well their solution performs.

*Ken Van Orman is a product manager for Spirent Communications.*

