



Enterprise Buyer's Guide to Layer 3 MPLS VPN Services

Peter J. Welcher

Introduction

This month brings the text version of another seminar, the one I presented at MPLScon in New York City in mid-May, 2005. I hope that putting some words to it will clarify the posted slides. In addition, after all the prep work for two presentations in two months, economizing on the effort of writing seems like a very good idea right now!

Due to space limitations, we'll focus on introductory material and MPLS L3 VPN's in this article. A following article will then discuss L2 VPN's, whether MPLS-based or not.

For those who'd rather read the presentation, it can be found at <http://www.netcraftsmen.net/welcher/seminars/mplscon05-buyersguide.pdf>. It is complementary to this article.

What are MPLS VPN Services?

In listening to all the presentations at MPLScon, it struck me that there's something rather basic that the MPLS, Service Provider, and journalistic communities have maybe not communicated very clearly to the networking public. To see this, realize that MPLS comes in two very distinct flavors:

1. **Internal MPLS** is where you already have links of various kinds, and you're deploying your own internal MPLS. This is usually done by Service Providers or Large Enterprise / Government organizations. Any use of VPN's is for private links or private routing tables, presumably to isolate and provide security between sub-entities (business units, government departments, etc.).
1. **Buyer of MPLS VPN Services** is where you're buying a WAN or MAN which just happens to be based on MPLS.

There's one other very solid reason I think this distinction is important. **If you're a buyer of MPLS VPN Services, you do NOT need to start by learning MPLS first!** All the MPLS is on the provider side. Knowing some MPLS can't hurt, but isn't a pre-requisite for getting started. You're buying a WAN / MAN service. Knowing what special things the provider can do to smoothly fit your environment may be more important. So if you've been looking at MPLS books, maybe they looked a bit complex, maybe you decided to wait another 6 months before thinking about buying MPLS VPN service ... well, you don't have to put it off any longer!

This article is intended for those in group 2, the buyers of MPLS VPN Services. It represents some of the things that I think you DO need to think about before ordering up some MPLS VPN Services. We'll indirectly touch on some non-MPLS LAN Services as well, in a later article. But that is not the main focus of this article. The focus is on understanding what's different about Layer 3 MPLS VPN, compared to traditional WAN services. And how those differences may affect you.

If you're in group 1, implementing your own MPLS, well, I'd love to talk to you as a consultant, but I'm not sure I have any article-length advice to offer. Internal MPLS design is something Ivan Pepelnjak, Jim Guichard, and others have written books about!

Why MPLS VPN Services?

I hear several motivations for buying MPLS VPN Services:

- 1 They're trendy
- 1 My boss told me to
- 1 Outsource all routing headaches
- 1 MPLS full-meshing is good for IP telephony
- 1 More WAN bandwidth for less

My personal take: the others may be plusses, but it's that last item that gets people looking. I'll stick my neck out: the benefit of full-meshing can be debated. It is a plus for MPLS VPN, but it seems to be getting more emphasis than it deserves. Delays over providers' high speed backbones are so low lately that, for example, regional back-hauling within the U.S. shouldn't matter. A delay of 30 msec one way is quite tolerable. Yes, with T1-ish Frame Relay or ATM, the delay can be noticeable, and meshing matters more there.


Terminology

The Provider Edge (PE) is the edge provider device, connected to your Customer Edge (CE) device. The PE is usually at the provider premises. If the provider manages the router at your site, it is a managed CE.

L2 Versus L3 VPN

Layer 2 VPN's provide a (relatively) clear data path at OSI model layer 2, a pipe down which you send bits. In the MPLS VPN world, they're quite likely to be handed off to you as Ethernet (FastEthernet, Gigabit Ethernet). The Ethernet may be point-to-point or multi-point. However, Frame Relay and ATM can also be emulated over MPLS. Non-MPLS L2 Ethernet services are available that are based on switching, usually using 802.1q "QinQ" tunneling. They may also be based on access devices multiplexing traffic directly onto SONET, in effect "new-wave TDM".

The reason I put in the word "relatively" above is that the transparency you get can vary by provider. Do they pass CDP (Cisco Discovery Protocol) and STP BPDUs (Spanning Tree hellos) between your devices?



Layer 3 VPN's use any form of connectivity to reach the provider, it doesn't really matter what (leased line, FR, ATM, Ethernet). Layer 3 VPN's commingle your routing securely with that of your provider. Your router doesn't peer with the one at the other end of the circuit, it peers with the PE router. That's novel to most of us. We'll look at some of the design implications below.

Depending on provider equipment vendor, the routing handoff between PE and CE can be based on almost any routing method: static, RIPv2, OSPF, EIGRP, eBGP, and ISIS. (But who'd want ISIS?)

Requirements: Questions to Ask Yourself

The first thing you really need to consider is whether you're looking to outsource, and how much. Do you want managed links? Managed routers? Managed routing? That's the first significant thing that's different about MPLS VPN Services, compared to leased lines, Frame Relay, or ATM: you may not be buying Layer 2 connectivity. L3 MPLS VPNs mix your routing with the Service Provider's routing in a secure, controlled manner. If you're a retailer or manufacturer with two overworked networking people, outsourcing routers and routing headaches might seem very attractive. If you're a large organization with solid in-house expertise, outsourcing routing might seem very unattractive. That's particularly so if you've been burned by previous exhibitions of inconsistency or mediocrity in Service Provider (SP) skill levels.

Layer 3 MPLS VPN can be provided with managed or unmanaged CE router. It is a routed service (hence L3). If you don't want to share routing with, or offload routing upon, your service provider, then you should perhaps think about a L2 VPN service.

A second question is: what is the local availability of L2 VPN services? Right now they're a bit spotty. Verizon is ramping up their TLS (Transparent LAN Services). Verizon TLS can be quite cost-effective and there are indications it may soon be much more widely available. I do have some reservations about relative risk and trusting such a service, depending on the underlying technology. We'll discuss this later on, in this or another article. The concern I have is not Verizon, it is Spanning-Tree Protocol. But you might decide the price is right and proceed anyway.

A third question to consider is: do you require a single SP or dual SP for redundancy (and negotiating leverage)? With a single SP, you can get locked in, particularly if you collocate gear at their site(s) -- the hassle of migrating to a new SP limits your ability to react to poor service. With dual SPs, your costs may be higher, but your network may stay up even when one provider is having problems. By experiencing two SPs, you can pressure them on pricing, and you can also compare the levels of service they actually provide. (Experience says that sales folks make SLA's sound great, but what you really get can be somewhat different!)


The remainder of this article concentrates on L3 MPLS VPN service.

Design Hint: Backdoor Routes

You have an existing link, say FR or ATM or IPsec VPN. You'd like to keep it as backup for your shiny new MPLS L3 VPN. You're fine if your CE gets internal OSPF or EIGRP routes, or eBGP from the PE. But if you have a managed CE, and if the provider insists on doing eBGP from PE to CE then redistributing into your IGP (OSPF or EIGRP), you get handed external routes, with very unattractive administrative distance. Your traffic will then use the older slower link rather than the new MPLS VPN. The same happens to you if you do the CE-based redistribution from eBGP to your IGP, and if your backup link router isn't the same as the eBGP-speaking CE router.

This isn't a show-stopper. You just need to think about the routing handoff and negotiate what you want, if you can get it.

Cisco routers do allow the PE router to provide internal OSPF or EIGRP routes to the CE router, despite having transported the routes using Multiprotocol BGP. But I haven't seen providers that are interested in providing an internal OSPF or EIGRP handoff to the CE. If you're buying the service, you probably want simple routing, so why would you want to have to speak BGP? Furthermore, if your MPLS VPN covers say 200 sites, do you want to have 200 eBGP-speaking CE routers? Or would you rather have them speaking OSPF or EIGRP?




This seems like a competitive business opportunity for the providers. The sales pitch is basically "we become your OSPF area 0" or "we become your EIGRP core". The drawback from the provider side is, there's lots of provider expertise or at least knowledge of BGP, less so for the internal gateway protocols (IGP's). Acquiring that expertise (experience!) is not inexpensive.

Design Hint: Dual Providers

If you buy managed services from two providers, they will each put a router at each of your sites. How does traffic get load balanced across the two routers? What do you do for failover? The issue is getting two different providers to run a common First Hop Routing Protocol (FHRP), that is HSRP, GLBP, or VRRP, with each other's router. It might be simpler to put a router or L3 switch at each site, but then you've created a single point of failure, equipment and cost, probably all the things you went to managed VPN's to get away from. Two site routers just makes it worse.

So it seems more likely that if you buy L3 services from two providers, the services will terminate on one or two CE routers that you control. For routing, you can run eBGP to each provider, and iBGP between the two CE routers (if you have two). If you can buy an internal OSPF or EIGRP handoff from the providers, make sure to buy the same handoff from both. It is highly inadvisable to get OSPF from one, EIGRP from the other, and do some form of redistribution at many sites. Route redistribution is great when used in limited and controlled manner. When you overdo it, it's just like beer. Too much and you end up with a big headache and an upset stomach.



Bypassing Provider Routing

Suppose you want to do your own routing, but all you can buy is L3 VPN service. Or perhaps HQ bought and imposed the MPLS VPN service on you, and your sub-organization wants to retain routing control. (Usually a losing political battle, but that's OSI layer 9). You can run GRE tunnels across the MPLS cloud, using the provider routing for connectivity between CE routers. And then run your IGP over the GRE tunnels. Note that this may entail a performance hit, and require dealing with MTU issues.

Other Questions for Your Provider

Under the hood, MPLS L3 VPNs use multiprotocol BGP between provider PE routers. You might ask some questions about this (good luck getting answers):

- 1 How robustly architected is the MP-BGP? Dual route reflectors or mesh of i-MBGP peers?
- 1 How fast is convergence? BGP timers? Scan timers in and out?
- 1 What IGP is the provider using in their core? Is it isolated from customer routes?

The convergence time is a real factor. Even if the edges look like OSPF or EIGRP to you, inside the provider it's MP-BGP. BGP is a lot of good things, but convergence is generally nowhere near IGP speeds. If that matters to you, you could test speed of internal route change propagation across the provider.

The part about provider core IGP is stability and convergence. If it carries customer routes, routing tables are huge, lowering stability and speed. If it is single large-area OSPF or ISIS, as is common, that is not great -- but probably something the provider isn't going to fix quickly, or just because you didn't like it.


Some other questions to ask:

- 1 Oversubscription and SLA's: what bandwidth are you guaranteed?

- 1 Security measures, especially for managed CE's? Would they detect configuration changes? Audit trail? Secure NOC? Will they put you in touch with someone who can provide some detail sufficient for due diligence questions about why should you believe the NOC is secure?

Since MPLS VPN routes your traffic directly across the provider MPLS network, any site can talk to any other site. Hence your CE links are multi-point links.

This has security impact: one central IDS at HQ doesn't see all the traffic anymore (unless you use a hub & spoke VPN design).



This also impacts QoS: if the provider doesn't provide QoS, you can't Do It Yourself, the way many of us have been doing for FR and ATM. The difference: with FR and ATM point-to-point links, you could prioritize your own traffic and allocate bandwidth, but not do anything about oversubscription with the FR or ATM cloud.

The problem with trying to do that for MPLS L3 VPN is that oversubscription by your own traffic may cause congestion on the egress link from the PE router to your CE router. And any QoS for the PE router will have to be done by the provider.

Conclusion: if you want QoS, it better be on your shopping list as something you buy from the provider.

Summary

Next month we'll look at L2 VPN's in general, not just MPLS-based L2 VPN's. The article will include some technical gotcha's and things to think about, based on my experiences to date.

Here are some links about MPLS VPNs.

Title	URL
For more on the underlying technology, see my article: <i>BGP and MPLS-Based VPNs</i>	http://www.netcraftsmen.net/welcher/papers/mplsvpn.html
AT&T IP Frame Relay Service (MPLS L3 VPN)	http://www.att.com/gov/ip_framerelay.html
Cisco page about Equant MPLS L3 VPN	http://www.cisco.com/en/US/netsol/ns341/ns396/ns172/networking_solutions_customer_profile0900aecd80129133.html
Verizon MPLS VPN	http://www22.verizon.com/enterprisesolutions/Includes/SiteUtilities/JCMSSkeleton.jsp?filePath=/Anonymous/Default/ProductDetail/Data/IPVPN_p.html

Your comments, questions, and suggestions for future articles are of course welcome! See below to decipher my email address.

Dr. Peter J. Welcher (CCIE #1773, CCSI #94014, CCIP) is a Senior Consultant with Chesapeake NetCraftsmen. NetCraftsmen is a high-end consulting firm and Cisco Premier Partner dedicated to quality consulting and knowledge transfer. NetCraftsmen has ten CCIE's, with expertise including large network high-availability routing/switching and design, VoIP, QoS, MPLS, IPSec VPN, wireless LAN and bridging, network management, security, IP multicast, and other areas. See <http://www.netcraftsmen.net> for more information about NetCraftsmen. Pete's links start at <http://www.netcraftsmen.net/welcher> . New articles will be posted under the Articles link. Questions, suggestions for articles, etc. can be sent to **[pjw <at> netcraftsmen <dot> net](mailto:pjw@netcraftsmen.net)** (formatted this way to fool email harvesting software).

6/13/2005

Copyright (C) 2005 Peter J. Welcher