

Is It Now Or Never For MPLS?

David Rohde

Enterprises now have the opportunity—and perhaps the obligation—to install the real deal without the illusion that it’s “just like frame relay.”

Haven’t moved from frame relay to MPLS yet? Cheer up! You can now skip some of the frustrations of the earliest migrants to Multi-Protocol Label Switching. Besides, you may need some cheering up to offset the stress that comes from the knowledge that a move off frame relay will, sooner or later, probably be mandatory.

Taken in by the carriers’ early pitches of MPLS as “sort of, kind of, just like, good old frame relay,” many of the early seekers of MPLS bids had a hard time nailing down deals where MPLS really could save them money, especially if they were focused on a set of connections that mimicked their frame relay network design.

Many enterprises now know better. Fortunately, carriers’ sales teams—some of them, anyway—also have learned not to pitch MPLS as just “one step up” from the last technology. For most carriers, MPLS is now for all intents and purposes a true IP virtual private network, linking authorized sites in a connectionless, any-to-any manner over a native IP backbone, with real, measurable distinctions among classes of service.

So enterprises now should buy MPLS that way, not with the mindset of frame relay network managers intent on securing perceived “reserved bandwidth” via permanent virtual circuits on predefined site-pairs to avoid any possibility of dropped packets. To show that they understand the difference, enterprises with a progressive mindset are finding it advantageous to go to market stating that they want a secure IP wide-area network with specific characteristics that meet their own application needs, rather than just call out for “MPLS” and let each carrier propose what they may.

Although they will certainly indicate that they want MPLS proposals prominently among the responses, these enterprises will give potential

suppliers a set of goal-oriented parameters regarding their desired network performance limits, maximum throughput, tiering of the network (flat or hubbed), and general inclination on whether they want managed or unmanaged service.

What savvy enterprises avoid saying for sure is whether they want voice to run over the network right off the bat, and whether they will or won’t consider other IP services that are alternatives to MPLS. If you make this mistake, some carriers will then immediately box you into a managed network solution and limit the way you can hub your network.

And while it may be anathema to many types of enterprises to consider this, the carriers also have available inexpensive feature enhancements to dedicated Internet connections that provide label-switching security and traffic engineering to your most critical connections without a full-blown MPLS installation.

The other side of the coin is that you can’t wait much longer to get going. Among the major U.S. carriers, only Sprint is explicitly telling customers that frame relay support won’t last through the next term of their contract. But AT&T among others is facing the loss of support of portions of their legacy frame relay/ATM equipment in the U.S. from Cisco, and it’s reasonable to assume that loss of all support will follow by the end of the decade. That’s why many enterprises using frame relay are hearing rhetoric in their discussions with account teams from AT&T, Verizon and other carriers that only an MPLS deal can get them the best rates. So the first step is to know what that really means.

Class Of Service—Know What You’re Buying

To understand today’s MPLS procurement picture, it helps to reorient network managers’ historical notion of “class of service” as merely a feature or option with a surcharge. It may have started out that way in MPLS, where original pricing schemes were essentially imitations of frame relay’s permanent virtual circuits. But now class of service generally *is* the pricing scheme once you have your access circuit and your port. The difficulty comes in comparing the numerous ways the carriers tie these CoS prices to the network.

David Rohde is a senior consultant with TechCaliber Consulting, LLC, a provider of telecom procurement, rate benchmarking and contract compliance services for national and global enterprises. He can be reached at drohde@techcaliber.com.

Carriers have many picturesque ways of naming their classes of service. Before signing on the dotted line, you should correlate whatever they call these classes to the actual definitions created by the Internet Engineering Task Force:

- Real-time CoS marked Expedited Forwarding EF (voice)
- Video CoS marked Assured Forwarding AF4x
- Data (burstable) CoS marked Assured Forwarding AF3x
- Data (burstable) CoS marked Assured Forwarding AF2x
- Best Effort marked with the supplier's applicable BE DSCP marking

In the generic CoS operational model used by most carriers, whatever they call the real-time class of service is assigned a fixed bandwidth without bursting, then the premium data class or classes and the best-effort class are assigned different slices, and the bursting headroom is shared among all the classes except the real-time CoS. However, how the carriers price the CoS traffic prioritization feature varies quite a bit from carrier to carrier, ranging from no cost for CoS at all, to paying for each different network connection based on a complex grid of combinations of CoS and bandwidth. Even AT&T, within its own shop, has several MPLS products—notably managed-router MPLS called Enhanced Virtual Private Network (EVPN) and unmanaged MPLS called IP-Enabled Frame Relay—which do not have the same pricing scheme for the CoS.

This is one reason why carriers have had quite a challenge defining network-to-network interconnections (NNIs) among their MPLS services. Instead of an international standards-based NNI, they have had to create custom NNIs with selected MPLS partners to provide additional geographic reach for their own customers.

A further difficulty is that the carrier MPLS points of presence or provider-edge routers may be a lot further away from certain of your sites than you might think. The large carriers have dramatically stepped up their number of MPLS nodes in the U.S., but they are still likely to slip their backhaul costs into their CoS pricing schemes, particularly if they try to give you a site-specific pricing list rather than a general price list offering the same price for the same combination of access line, port and CoS.

Making each bidder pull away from site-specific pricing into an apples-to-apples comparison for equivalent speeds and classes can help uncover who can serve particular locations economically and who can't (see below for more on site-specific pricing). Carriers should be able to provide the same combination of bandwidth and characteristics like CoS at any site that asks for it; when they revert to site-specific pricing, it's often because their points of presence (POPs) are too spread out, or they are just not experienced enough in the service.

Convergence: Is It The Horse Or The Cart?

Perhaps because they can't think of anything else to say, sometimes carrier officials will say that the purpose of the wholesale migration to MPLS is to achieve convergence. They can say that, but don't buy into it.

The fact is, a doctrinaire reading of MPLS as your way into VOIP is likely to reduce your flexibility rather than increase it, for several reasons. A paradox of MPLS is that you almost certainly cannot use VOIP off the bat to make the network pay off, but once the network is in place and saving money on your critical data applications, you can add VOIP to save a lot more money. Don't let them put the cart before the horse.


There are several reasons why this is so. In the initial move to VOIP, the cost savings are almost always diluted by resiliency requirements, which limit access savings. And the real-time CoS and CPE equipment upgrades required to implement VOIP may dilute cost savings. The innards of individual carriers' MPLS traffic engineering policies may reveal them to be as skittish about voice traffic congesting their IP/MPLS networks as they would be on any other network, whether they call it an IP-VPN or not. Much depends on the codec used and how many calls they are really willing to fit onto a 1.5-Mbps MPLS port.

Even to be allowed to use a real-time CoS, some (though not all) carriers may force you to consider a managed service, since they will not offer a suitable jitter guarantee without your doing so. Thus, getting married to the idea of using your MPLS network for convergence may dramatically limit your flexibility and leverage in your initial MPLS procurement, constraining your ability to consider a decent range of service provider options.

And don't worry about your time frame to get to VOIP over MPLS. The carriers may be laying plans to turn down their frame relay networks, but they are not planning to shut down circuit-switched voice—not this decade, anyway. Rather, your principal timing variable may be where you stand on upgrading your voice CPE across your enterprise, since IP-telephony is now the standard there rather than just the "next generation."

But thinking through how you're going to *ultimately* handle voice also will have an impact on your MPLS network design. In pan-European and pan-Asian networks, enterprises and their carriers have a tendency to create hubs for their branch offices in each country. Hubbing, whether in-country or in-region, opens up new possibilities in the form of some of the increasingly available MPLS custom NNIs among carriers. But it also tends to introduce unacceptable latency to the voice-grade class of service.

What about off the network? With some carriers, off-net voice usage may be priced in bundles of minutes (rather than explicitly toll-based) but is essentially still usage-based in the sense that it



Your MPLS decision-making should reflect, not drive, your VOIP migration strategy

The carriers may need to know more about your applications and performance expectations than they used to

involves a fixed monthly price based on the number of concurrent voice calls that is pre-defined and provisioned by the enterprise. There are usually no usage charges for on-net traffic, but assuming an enterprise chooses an MPLS-with-VOIP service that is carrier-managed, additional recurring charges will apply. There are likely to be several of these—charges for the rental and maintenance of each router, a separate management charge for the management and monitoring of the router and MPLS connection, and various charges for Moves, Adds, Changes and Disconnections (MACDs).

And remember, moving to a VOIP scheme could lead to shortfalls in current contractual annual dollar commitments. The move could also undermine obligations in some carrier contracts to maintain the traffic mix in particular proportions established at the beginning of the term—a nefarious fine-print practice called “monitoring conditions”—and could also violate some regulatory restrictions in foreign countries if not carefully checked out in advance.

Difficulties In Pricing The Service

Additional pricing gotchas abound in MPLS procurement. Some are even making encore appearances from the early days of frame relay.

Example: Carriers may charge separately for domestic vs. international classes of service, just as they originally charged up to four times as much for a frame relay port if there was an international PVC attached to it. In MPLS, there’s not much logic to this—the user’s connection is only to the edge of the network regardless of where the traffic ultimately is going. As a matter of fact, one of the reasons that MPLS first took off in Europe, and many European enterprises “skipped a generation” (to American eyes) by going from private lines to MPLS, is that international frame relay PVCs (which do cross country boundaries) were expensive in a way that MPLS network-edge connections don’t need to be.

Also, don’t fall into the trap of letting your selected carrier provide site-specific pricing, wherein each site has a specific price on a custom price schedule. What may look enticing and specific when you begin an MPLS deployment will leave you vulnerable as soon as you expand or change the configuration at an existing site. The resulting contract may have no constraint on the carrier charging whatever it wants on a new site, which it can always claim has some different characteristic (by virtue of distance to the POP, unique port or CoS bandwidth, or anything else) than an ostensibly similar-sized location.

The carrier may even revert each new site to “Service Guide” pricing—the set of unrealistic, way-out-of-market prices that are buried deep in the carrier’s website. So make sure you get a pricing schedule that can apply to both current and future sites in a predictable way.

MPLS Network Design vs. Frame Relay Network Design

MPLS really is an IP-VPN, which means that ATM, frame relay and private line topologies cannot be directly translated into an MPLS network, for several reasons. For some users, you might actually be able to reduce hub site bandwidth requirements vs. those of a frame relay network, because of the more dynamic nature of any-to-any connections.

But at the same time, bandwidth requirements may need to be broken down by CoS, and CoS throughput bandwidth definitely will need to be included at hub and remote sites, taking into consideration application traffic utilizations and traffic flows. That’s particularly true for enterprises that currently employ partial or full frame relay PVC meshing, where the enterprise has multiple server farms/datacenters or requires traffic to reach anywhere because of any-to-any applications such as VOIP.

Thus, the carriers may need to know more about your applications and performance expectations than they have in past procurements of new networks. And the more you can define those according to the IP/MPLS specs they can recognize, the better.

So you need to investigate with each provider the MPLS network’s physical and logical topology; any Layer 2 access used; whether the MPLS traffic is transported over the providers’ shared router public Internet IP or non-shared IP infrastructure; and their MPLS network bursting policies, which may be far different (and potentially more expansive) than their traditional frame relay policies.

Additional questions that typically arise include any enhanced MPLS services that are available, such as network-based firewalls or Internet datacenters, and the routing protocols supported on your customer premise equipment. Hint: Be prepared to implement eBGP on your customer premise routers!

All of this will affect your ultimate service level agreement with your chosen MPLS carrier. With MPLS, it’s easier than ever for carriers to slip into the habit of only tracking their performance within the carrier cloud, since MPLS label tagging is performed at the provider edge (PE) router rather than the customer edge (CE) router. SLAs should include site-to-site, and not just PE-to-PE router metrics, and should be based on site availability, not “network” availability.

Here’s where styling your procurement project as not strictly MPLS can help. In some regions of the world, MPLS still is not either the most economical nor most popular choice. In Canada, long-distance, point-to-point Ethernet services have really taken off. And in China, historically, nothing has been cheaper than private lines. The reality, however, is that enterprises want to create pan-Pacific regional MPLS networks, and many

carriers are creating MPLS NNIs with one of the two principal Chinese carriers. Letting the carriers know you are ultimately focused on the end-goal of an efficient IP network, not a specific carrier solution, will help sharpen up their responses to you.

Surcharges Will Be An Issue

As extra motivation for a move off of frame relay, a pleasant and unexpected financial surprise may be in store for *some* frame relay users when they move to MPLS. Because frame relay is considered a “basic telecommunications service” by regulators, the carriers are supposed to levy universal service surcharges on the service (or, more precisely, they are supposed to count it in their own universal service obligations, which they pass on to users). But MPLS is an IP service that is not under the same rule, and such surcharges should not appear.

Paradoxically, for many larger frame relay users the same *comparative* advantage won't be there, because AT&T long ago quietly removed such surcharges on networks that employed frame-to-ATM interworking (FRASI), on the logic that ATM was never classified as “basic telecommunications” and all such FRASI links have ATM on at least one end. So the move to MPLS for users with this common set-up might not include that comparative advantage.


And many factors suggest that telecom regula-

tors won't allow the entire universe of enterprise networks to move completely off the surcharge regime. Certain configurations of VOIP are already being targeted for circuit-switched-type universal-service surcharges, and while it's not clear that voice-over-MPLS falls into this category, the issue bears close watching.

Conclusion

It's now more true than ever that telecom managers need to be facilitators and organizers both up and down their own organization charts. The move to MPLS can help in this process. In the complex world of IP voice and data, there's a premium to being able to present to current and potential service providers your ideal network and what it has to accomplish in terms of applications, and at what kind of price. That means getting your own senior managers on board a plan, so that the service providers don't wind up presenting you with options that move you to a next-generation network without actually saving you any money, or even wasting bandwidth.

The reality of a looming change in data-networking generations has a way of concentrating minds. Each enterprise manager has the power to decide whether it will be a headache or a boon to their own competitive position *vis-à-vis* the few remaining competitive carriers across the nation and around the world□



Regulatory effects may end up saving you money