What Price A New Network?

David Lee and David Rohde

Don't let marketplace anomalies trip you up on the way to your preferred architecture.

Most eras in corporate networking tend to feature one technology that everyone talks about, and another one they actually employ. For example, users have been kicking the tires for years on successive waves of IP virtual private network (IP-VPN) carrier services, and even installing their own CPE-based IP-VPNs for many applications. But until recently they almost always defaulted in their carrier contracts back to frame relay at branch sites, often combined with ATM (asynchronous transfer mode) at headquarters and datacenters.

Those days are going fast. Suddenly, there actually is no single, dominant technology for enterprise wide area networking. Competitive procurements on the street routinely demand side-by-side quotes for some combination of frame relay service, RFC-2547bis compliant Multi-Protocol Label Switching (MPLS) network offerings, managed-router services to accompany the transport, and various offerings built around carriers' access links to the public Internet.

What's more, some competitive procurements now also include solicitations for national virtual LAN services built around plug-and-play Ethernet connections or compliant to the IETF draft for Virtual Private LAN Services (VPLS). Alternatives, with credible footprints, are even being offered to the standard T1 for mission-critical corporate networking services (and not just generic Internet access). Middle-market and retail industry users in particular will consider using services like DSL and broadband cable to bring a significant number of sites onto the network, opening the door for a new generation of broadband access aggregators to get into the enterprise networking arena on the services side.

The extra choice is a great boon to users, isn't it? Mostly it is. The willingness of users to move on from frame relay naturally makes it easier for them to consider competitive alternatives to their incumbent carrier, and that's always a boost to negotiating leverage. But—and you knew there was a but!—the current environment also presents a series of challenges:

- The new services tend to present pricing anomalies where many users least expect them.
- Account teams within individual carriers tend to have trouble presenting the terms and service features of each new technology in a consistent way.
- Corporate network managers tend to hang onto old network sizing assumptions and network management habits from previous wide area network technologies, needlessly failing to take advantage of certain cost efficiencies inherent in what they've bought.
- Most ominously, the predominant service providers for these multiple technologies are getting bought, and their purchasers have promised operational and billing savings to Wall Street—all of which suggests that the carriers ultimately will have to choose which network platforms are going to stay and which are going to go.

No one wants to be caught in a guessing game of whether to bank on the tried and true or the fresh and now, so here is our advice to navigate the waters of today's technology choices at procurement time.

Carriers May Soon Force The Issue

The very fact that there are several predominant platforms and backbones rather than one raises the uncomfortable question of which platform will survive. In addressing the question, don't just focus on the companies that are undergoing mergers—AT&T and MCI. All long distance carriers and RBOCs are in the process of rationalizing their networks.

After all, carriers need to take their cue from their own suppliers, where upgrades to frame relay and ATM gear are coming to the end of their road. Given that some carriers standardized on Cisco's BPX 8600 series of multiservice switches

David Lee and David Rohde are consultants with TechCaliber Consulting, LLC, a provider of telecom procurement, rate benchmarking and contract compliance services for national and global enterprises. They can be reached, respectively, at dlee@techcaliber.com and drohde@techcaliber.com.

for frame relay and ATM access, and Cisco has declared this platform "End of Sale," these same carriers need to plan their migrations away from this equipment and onto a new platform.

Early on, carriers needed to structure their MPLS efforts around their existing ATM backbones. But now one significant wireline carrier is in the process of consolidating its MPLS transport services onto its IP backbone, so that its enterprise MPLS traffic and public Internet traffic ride the same core. Enterprise and Internet traffic will be separated and controlled via separate provider-edge (PE) routers and through MPLS Route Distinguisher/Router-Target (RD/RT) addressing. This theme of a common IP backbone in conjunction with traffic prioritization via MPLS or VPLS to enhance the user experience—and potentially charge a premium for mission-critical bits versus Web surfing bits—was presented by a number of speakers during last spring's MPLScon 2005 conference in New York.

Frame/ATM switches are nearing "end of sale"—will the associated services follow?

Some long distance carriers, strained for capex dollars and at risk of selling off once-prized wireline assets that could ultimately even include frame relay, may be making transitions in an even more aggressive way. If you don't intend to migrate technologies in your next contract, watch for the appearance of contract clauses from the carriers giving them the right to "de-select" or force migration from older fast-packet technologies even within the next term of your contract.

For global networks, where MPLS is closer to being the standard rather than just an up-and-coming technology, carriers are already saying that frame relay is of little interest. Some global carriers have been telling users in procurements that they might as well forget about frame relay, since the carrier would vastly prefer to bid MPLS.

Sizing The Network Properly

That may be why you don't hear as many presentations these days making a clear linkage between frame relay and MPLS, or suggesting that there's little involved in making a natural upgrade. In fact, it can be easy to lose sight of what you truly need to do to actually obtain real cost savings from any conversion.

Many users recognize that while they may not have to swap out their routers, they have to do a lot of coordination with regard to potential Cisco software or memory upgrades and routing protocol support. For example, many enterprises continue to employ the Cisco-proprietary routing protocol Enhanced Interior Gateway Routing Protocol (EIGRP) in frame relay and other data services. But some leading MPLS providers don't support EIGRP over their MPLS-based services, and require that users employ the Border Gateway Protocol (BGP) or resort to static routing.

In an attempt to gain advantage, competing MPLS service providers may preserve EIGRP routes across their MPLS networks. But check the fine print. You might have to get an "approved" contract amendment to achieve this, and this support may carry an extra monthly recurring charge per site. Or EIGRP support may be promised by account teams but not be fully supported—something you'd rather not find out during negotiations or, worse, after contract signing.

An even more surprising change may have to do with basic carrier network engineering issues. A generation of AT&T frame relay users in particular has grown up with the notion that fast packet networks should be established under tightly prescribed and proportional parameters. Because of AT&T's "ReliaBurst" (or closed-loop) congestion management scheme, users were told to buy a level of committed information rate (CIR) on their permanent virtual circuits (PVCs) that would ensure them competitive bursting rights and guard against dropped frames. Many experts believe that notion is now outdated—that AT&T's frame relay network has enough capacity to routinely accommodate bursts up to the port speed.

But under AT&T's MPLS offerings, the network sizing issues take on a whole new twist. AT&T generally has two MPLS service choices: IP-Enabled Frame Relay (for MPLS transport only) and EVPN (managed-router MPLS service). Each uses what are known as EPVCs that simply deliver the traffic to the nearest PE router rather than carry it across the network cloud. Many initial users have found they can provision a lower-bandwidth EPVC, particularly when they only require one class of service per location. In other words, they see no reason to let their EPVCs turn into a revenue enhancement tool for their carriers.

In these kinds of MPLS networks, the purpose of the virtual circuit is not to carry traffic across the cloud end-to-end with the risk of discard eligibility, as in frame relay. Instead, the EPVC brings traffic only to the edge router associated with the customer site. In AT&T's IP-Enabled Frame Relay service structure, the EPVC acts as a mechanism for apportioning classes of service, with progressive charges—the more EPVC bandwidth, the greater the monthly recurring charge.

For EVPN, class of service charges are incorporated as EPVC charges, and they vary by band-

It can't hurt to give alternate carriers a fresh look

width and priority (higher priority carries a greater charge). And yet, except for Class of Service 1 (COS1), traffic prioritization is done only at the egress (destination) PE router.

In this scenario, users could easily make two cost-busting mistakes:

- They might subscribe to more EPVC bandwidth than they need, thinking that frame relay style "bursting" needs to be protected across the network, when in fact, with the glut of long-haul fiber, the core is kept at a low utilization and features no traffic engineering. The bottlenecks occur in the egress PE routers, but with only the enterprise's MPLS port bandwidth, not any EPVC subscription mechanism, gating the performance.
- They might subscribe to too high a class of service. For many initial MPLS users who have simply moved over from frame/ATM but haven't yet implemented two or more classes of service, Class 3 out of a total of 4 classes generally serves as the frame relay-equivalent class of service.

Complicating the matter is the fact that other carriers handle the bandwidth and class of service charges differently. MCI has five classes, including a new class, just below real-time, for video or very low-latency data. Charges for the non-real time classes of service (the four data classes) are "free" in that multiple classes can be provisioned at various speeds without incurring progressively higher monthly recurring charges. The real-time class of service, termed Expedited Forwarding (EF) or Gold, does have a monthly recurring charge proportionate to the bandwidth provisioned, as this is protected traffic categorization for voice over IP (VOIP).

Sprint doesn't charge for any of its four classes of service, including its real-time class, but does limit or caution users who have a very high percentage of real-time bandwidth.

So in general, large committed bandwidth on the MPLS EPVCs is necessary only for parceling out traffic into different queues for different classes of service. Some users have even implemented up to 45-Mbps MPLS ports and yet chosen only 16 kbps of EPVC bandwidth. This reminds many people of the classic Sprint approach to frame relay, with 0 committed information rate on the frame relay PVCs. But the difference is that in MPLS, assuming the destination port bandwidth is sufficient to support offered traffic loads, the excess traffic isn't dropped. It goes into lower priority queues or the best-effort class of service queue, only competing with other traffic within the enterprise's VPN at the egress PE router.

Big And Little Bandwidth—Both At The Right Price

The increasing use of very high bandwidth connections presents some other unique pricing issues. In recent years, users of frame/ATM interworking branch-to-host services have become used to very significant price-downs in the cost of T3 ATM ports and related connections, to the point where a T3 could cost only 3–4 times a T1 port on the same network. But when they've explored other, cheaper network alternatives, such as dedicated Internet access, customers have often found that their T1 costs come way down but their T3 connections remain surprisingly high in price, with nowhere near the price compression across bandwidth seen in the frame relay/ATM combo.

One reason that Internet cost per Mbps doesn't track with traditional frame relay and ATM service relates to the cost of peering with other ISPs. Unlike closed frame and ATM networks, ISPs need to treat transit traffic (destined for other networks) equally as well as its own on-net packets, yet there is no price differentiation for transit versus on-net traffic.

All this helps explain why metro and national Ethernet connections, provisioned as part of an MPLS implementation, are catching on. In services known generically as Virtual Private LAN Service (VPLS), enterprises can implement 100-Mbps or 1-Gbps Ethernet local access with various port bandwidths up to the local access bandwidth limit, all supported with traffic prioritization (see *BCR*, June 2005, pp. 40–49).

Service providers such as Masergy have coupled VPLS and MPLS services with flexible operations, administration and management (OA&M) systems that allow a customer to "dial-up" or "dial-down" their port speed without the need to send a fax or email to the service provider. That's driving innovation into the traditional carrier community.

Unlike the original generation of "transparent LAN services" provisioned within metro areas, VPLS can provide true class of service capability even on a bridged, rather than switched, connection among key sites.

At the other end of the spectrum, the major carriers subcontract broadband aggregation to various DSL and cable broadband partners if a customer deems broadband Internet access suitable for their business model. Customers generally don't fully trust the reliability of DSL and cable broadband access compared to the tried and tested leased line.

However, in the case of both Ethernet and broadband services, one of the biggest dilemmas in the transition period is whether to deal with the same old carriers, or find value in some new carrier that your boss has never heard of. It can't hurt to include one of the new generation of broadband access aggregators in a competitive RFP, so long as you socialize the option throughout your organization rather than expending negotiating time with a vendor who has no chance of winning the business.

Preparing To Run Voice

Enterprises considering any VOIP transport over their WANs must consider the incremental costs

for real-time quality of service (QOS) support, the incremental MPLS port and real-time class of service bandwidth, and the cost to packetize the TDM voice from an existing PBX.

The one area where true QOS does make a difference is in support of VOIP both in the LAN and the WAN. Enterprises must ensure that QOS is maintained desk-to-desk and across the entire MPLS cloud. Multinational enterprises need to investigate whether their carrier's global MPLS solution is implemented with multiple BGP Autonomous Systems (AS'es). If so, investigate with the prospective carrier how QOS is maintained between border AS routers to ensure a low-latency, QOS-transparent transition. In addition, not all carriers have a ubiquitous global MPLS footprint, and thus have to employ MPLS Network-to-Network Interface (NNI) arrangements with other MPLS carriers to fill in gaps. Make sure to find out which provider is actually delivering the service in your geographic region of interest.

Incremental VOIP traffic needs to be accommodated by the corresponding incremental MPLS port bandwidth and real-time class of service. Since real-time traffic must be protected to ensure low latency, low jitter and low packet loss, MPLS providers discard, at the ingress router, any real-time traffic that exceeds the provisioned class of service bandwidth.

Consequently, the real-time class of service bandwidth must be sized to accommodate the total number of concurrent VOIP calls at a specific site. These incremental local access, MPLS port and class of service monthly recurring charges need to be considered in any analysis of cost savings compared to the declining PSTN marketplace. Finally, an integrated access device (IAD) such as a router, and its corresponding costs, needs to be deployed to packetize the PBX's TDM voice prior to transport over the MPLS network.

Hidden Costs—And Opportunities—On The Surcharges

The move to any kind of new corporate data network always carries one extra adventure with it—the mystery of whether the heavy universal service surcharges and related fees apply.

The root of the mystery is that regulators and the carriers define the dividing line between basic (or regulated) and "enhanced" (or nominally unregulated) services at different places. Traditionally, traffic that could be defined as "Internet" would escape regulation and surcharges, while basic services would fall under the regulation/surcharge scheme. Certain rulings in the mid-1990s specifically brought frame relay under the surcharge scheme, whereas its status had been unclear before.

But in recent years, AT&T often has exempted frame relay traffic from these surcharges if it's provisioned under an "interworking" scheme with unregulated ATM—a practice not always followed by its competitors. This has sometimes given AT&T frame/ATM customers a significant break compared with the carrier's competitors, given that the universal service surcharge alone currently stands at 11.1 percent.

Industry experts generally assume that MPLS traffic likewise can fall outside the surcharge regime. But since much MPLS traffic is defined by its service as "IP-enabled frame relay," surcharges could pop up on your bill anyway. And salespeople and account reps are notoriously inconsistent in their answers to the MPLS surcharge question, even within the same company.

Be sure you understand whether universal service and other surcharges apply to a new service

As a result, users put themselves in a much better position if they explain in their MPLS and other next-generation procurements that the surcharge question is open for discussion, and they explicitly ask, with regard to each rate element, whether universal service applies.

Conclusion

It's never "obvious" that one technology will save money over another until a contract has been signed and the bill is rendered. The now-rapid shift to new IP-oriented wide area networks, especially those wrapped in the mantle of Multi-Protocol Label Switching, provides both an opportunity for savings and a cost trap for the unwary.

Old assumptions about the nature of "virtual private networks" (in the broadest sense of the term) running over shared carrier infrastructures need to be challenged to make the transition pay off. Make sure the way you buy bandwidth fits both the technology you're riding on and the applications you're planning to use□

Companies Mentioned In This Article

AT&T (www.att.com)
Cisco (www.cisco.com)
Masergy (www.masergy.com)
MCI (www.mci.com)
Sprint (www.sprint.com)