

# How To Negotiate An MPLS Contract

Ben Fox and David Rohde

## Base your opportunity for savings on today's marketplace, not yesterday's frame relay contract.

In an ideal world, negotiating a contract for a service based on Multiprotocol Label Switching (MPLS) would simply be a matter of plugging a few different terms into a contract for a frame relay service, haggling over the price, and signing the deal.

After all, MPLS is said to be leading the charge for dedicated corporate IP virtual private networks (VPNs) because of its relatively direct upgrade path from widely-installed U.S. frame relay networks.

But corporate telecommunications professionals don't live in an ideal world. MPLS service can lead to better performance, more network flexibility, the opportunity for convergence and the chance for significant savings. But the process to get there involves a larger number of variables than users are accustomed to dealing with, even on their existing data network services.

For example, MPLS is often sold as a managed service by default, so merely considering MPLS introduces many users for the first time to a new array of charges and potential disagreements about what is and isn't in scope of the carrier's responsibilities.

And in MPLS, compared to frame relay, the unique roles in the pricing structure played by bandwidth reservation mechanisms offer fertile ground for carriers to confuse users. The carriers can roll out bugaboos such as "future savings" or "reduced total cost of ownership" rather than make hard savings part of the original bargain. Or they might introduce class-of-service (CoS) pricing that bears little resemblance to other carriers' CoS price structures.

Finally, it's usually not prudent to install a big MPLS network right as you're ripping out your

frame relay network. Migration costs—including the likelihood of running two networks side by side for a period of time—are thus a real consideration that must be paid for in order for the MPLS value proposition to pan out.

### Don't Tip Your Hand

Certainly now is a good time to consider MPLS. An MPLS service is a type of IP-VPN in which all sites can communicate with all others without needing a virtual circuit predefined in software for every site-pair, as is required in frame relay. In MPLS, each of your locations subscribes to a port connection at the edge of the carrier's network, where IP traffic is tagged with a label that signifies the closed user group defined by your VPN.

This label can be recognized and swapped throughout the carrier network regardless of whether the carrier's backbone is pure IP or a grid of ATM switches—thus the "multiprotocol" nature of the technology despite its intended use for IP applications. Because carriers such as AT&T and MCI have retained a frame relay user-to-network interface for their U.S. MPLS services, MPLS is proving attractive to users who wish to minimize the impact on their CPE and retain the security of Layer 2 fast-packet services.

So should you simply announce to your current carrier and prospective competitors that you're planning to move to MPLS? That may be the biggest trap of all, simply because of the current rapid de-escalation in costs for data networks in general (see *BCR*, March 2004, pp. 22–24).

In a declining rate environment, carriers love nothing more than having users tip their hand this way. Asking for MPLS rather than bids for both MPLS and frame relay is an open invitation for carriers to escape from the normal market-rate discovery process. If you don't know what the current market rates are for frame relay ports and virtual circuits, your incumbent probably will bid the MPLS network as if the benchmark were your existing, years-old frame relay rates.

The way MPLS is typically deployed in the U.S. makes this effect particularly pernicious. In

*Ben Fox and David Rohde are consultants with TechCaliber Consulting, LLC, a global consulting firm specializing in benchmarking, procurement and contract compliance on behalf of business users for their voice and data networks.*

the popular imagination, MPLS and other kinds of IP-VPNs rely on a single port connection into the carrier cloud, and a connectionless topology throughout the cloud to provide any-to-any connectivity. But services such as AT&T's Private IP VPN Service and Enhanced VPN Private IP (EVPN), as well as MCI's Private IP, do carry a charge that represents the MPLS bandwidth from the carrier edge into the MPLS cloud—think of it as a substitute for the Committed Information Rate in frame relay PVCs.

Especially at lower data speeds, this “e-PVC” with a Committed Data Rate or Committed Access Rate into the carrier cloud can be priced at something of a premium to typical frame relay PVCs, and a main goal of many negotiations is to knock this premium out. But unbeknownst to many users, the market prices of plain-vanilla frame PVCs are plummeting after years of stagnation. So failing to benchmark current PVC prices by CIR vs. the equivalent CDR or CAR rates in MPLS could lead you to overpay even when you think you aren't.

Another key pricing factor: class-of-service markings. Up until now, the CoS capabilities of packet networks have been pretty much a washout. Users have largely ignored various marketing attempts at frame relay class-of-service offerings. And even in ATM networks, where fixed cell lengths make quality of service distinctions inherent, users have largely stuck to the basic variable-bit-rate non-real-time (VBR-nrt) class of service for data applications, occasionally stepping up to constant-bit-rate (CBR) service for video traffic, but ignoring other variations.

For many prospective MPLS users, any-to-any connectivity is still a bigger technical driver than class of service, which they may simply find too difficult to take advantage of. But other users have firm plans to put CoS capabilities to the full test. For example, as frame relay networks have filled up, some users have found that transoceanic frame relay links have slowed down, and they plan to use MPLS's lean structure in the cloud—where intermediate switches just swap labels and don't have to worry about IP addresses and payload characteristics—to run real-time applications off a host computer on another continent.

Other users plan to run voice and streaming audio and video over their global MPLS networks by taking advantage of the real-time classes of service that the carriers have specifically designed for such demanding traffic types. Or they want to take advantage of MPLS CoS characteristics to give priority, and specific performance guarantees, to their global SAP traffic while throttling back non-business Internet traffic.

### Lining Up The Comparisons

But in their negotiations for these services, users are discovering that many carriers have relatively undeveloped service level agreements (SLAs) for certain classes of service.

Users should demand that guarantees for Site Availability, End-to-End Network Availability, Maximum Time to Restore, Network Transit Delay, Packet Delivery Ratio, Jitter and Provisioning are all included in an SLA for MPLS services. And they should make sure the carrier's SLA for real-time classes of service—sold at a premium to lower classes of service for standard data—are stringent enough to support the type of

voice and video that these classes of service are purportedly designed for.

And check the fine print. Carriers have been known to package their MPLS product offerings in such a way that the offer excludes the SLAs the customer really needs. For example, carriers may quote a maximum level for jitter (a critical para-

meter for the successful use of VOIP) but only include it in the SLA if the user purchases the carrier's own managed VOIP feature for the MPLS service. That won't help the many users who want to manage VOIP in-house even while the traffic traverses the carrier's MPLS cloud under the proper class of service.

Breaking the deal down into its rate elements is critical to sorting through these challenges. A common (though by no means ubiquitous) MPLS pricing structure is:

- **Access charge**—The charge for providing the local access circuit between the customer's site and the service provider's point of presence (PoP).
- **MPLS port charge**—The charge for the MPLS port required at the service provider's PoP
- **MPLS bandwidth charge** (also referred to as subscription bandwidth charge, throughput charge, ePVC charge, Committed Access Rate, or Committed Data Rate)—The charge for the MPLS bandwidth that the site being connected will require into the MPLS cloud.

Typically, though not always, the bandwidth charge will be assessed on a per-kbps basis (or suitable multiples thereof—the same as buying a frame relay PVC), and the charge will vary depending on which CoS is purchased. Then separate MPLS bandwidth can be bought for each class of service that the user requires.

But there are variations. In their home markets—say, AT&T in the U.S. and BT in the UK—the carriers may bundle the access circuit charge into the port charge. Sometimes a carrier may not have a separate MPLS bandwidth charge and instead levy different port charges depending on the CoS mix provided. Other times the “best

**Many carriers have  
relatively undeveloped  
SLAs for MPLS**

effort” CoS is bundled with the port, and all other classes of service are extra. And some carriers vary the bandwidth charge depending on whether the user needs to pass packets internationally or just within the same country where the traffic originates.

### Different Around The World

None of these pricing alternatives are necessarily better or worse, though different pricing structures will tend to suit different users. The carriers seem unlikely to normalize on a single pricing structure, and this diversity in MPLS pricing structures likely will be with us for the foreseeable future.

But some rules of thumb can help you sort through the maze:

■ For global networks, the savings from migrating to MPLS will likely be higher in Europe and Asia than in North America. Even using pan-European carriers, frame relay PVCs can be very expensive as soon as you cross national borders.

The same doesn’t apply to MPLS e-PVCs because instead of traversing the cloud, they cover the path into the network edge.

■ Watch for creeping installation charges. Many users have grown accustomed to knocking out all installation charges in their voice and data contracts. But with MPLS, you are often marrying yourself to the carrier’s budget for capital expenditures. When financially strapped carriers tell Wall Street that they are engaging in “success-based” capital expenditures, what they’re really saying is: “We’re going to hit up our customers for the costs of installing things like MPLS.” Your order for MPLS service may be the one that causes your carrier to go back to Cisco for the necessary cards and other upgrades.

■ MPLS connections come in a wide variety of bandwidths—anywhere from 64 kbps to 45 Mbps. Users who have previously committed to frame relay-to-ATM interworking will also need to replace their ATM connections with new MPLS connections. Since ATM T1, N × T1 and T3 ports have come down radically in price, make sure you don’t sacrifice those savings to MPLS without first comparing ATM bids or benchmarks.

■ It can be hard to immediately make the MPLS business case add up if you are replacing a domestic hub-and-spoke network, are only buying transport, and cannot quantify additional savings from related areas such as any-to-any connectivity. But if you’re negotiating a new deal now and anticipate moving to MPLS in mid-contract because of evolving business or network topology needs, don’t make a minimum annual commitment to your current data network (or your telecom contract as a whole) that’s too high. The very savings

you might achieve via MPLS any-to-any meshing in mid-term could put you into shortfall on your volume commitment, destroying any further negotiating leverage with your carrier.

■ MPLS can be a cost-effective way of buying managed services because the suppliers bundle the management and router in with the transport. Many users have seen their managed frame relay costs soar because the equipment maintenance and service monitoring come from a different product house inside the carrier than the transport piece. At least with MPLS, many carriers will take the approach of including management together with the negotiation of transport rate elements

because it’s inherently part of the service. That could potentially lead to a much bigger discount.

■ The any-to-any connectivity capability of MPLS often means that hub-site infrastructure can be removed or scaled down because less traffic actually traverses the hub site. Make sure to scale your hub-site (or

datacenter) connection into the MPLS network according to its actual traffic patterns, not according to typical frame relay calculations. Also remember that MPLS gives you the capability to host applications throughout your network, not just at the central site.

### Conclusion

As if this weren’t enough, suppliers have a bad habit of quoting individual prices for MPLS on a site-by-site basis as opposed to giving structured list pricing with associated discounts (as frame relay users are accustomed to). This makes pricing very un-transparent and means the pricing for site upgrades is not included in the contract, and hence every site upgrade or downgrade becomes another negotiation.

All the more reason to consider MPLS in a broad context, together with other data network proposals and, if possible, even your national or global voice spend. If the carriers see that you are intent on picking the best, most cost-effective network for your particular needs—not just to fit their migration strategy—you’re more likely to get the best possible deal □

**Watch for creeping installation charges and new traffic patterns with MPLS**

### Companies Mentioned In This Article

AT&T ([www.att.com](http://www.att.com))

BT ([www.bt.com](http://www.bt.com))

Cisco ([www.cisco.com](http://www.cisco.com))

MCI ([www.mci.com](http://www.mci.com))