Comparing Apples To Marshmallows And Other Tales Of MPLS

Editor's note: *TechCaliber Consulting Managing Director Ben Fox and Senior Consultant David Lee both contributed greatly to this column.*

The focus of Customers' Corner is typically commerce and policy—carrier contract machinations, or key regulatory issues that affect the user bottom line. That said, once in a while we venture into technical subjects, if only to burnish our geek *bona fides*. This is one of those times.

How do you evaluate bids when products are almost designed not to be comparable?

In the good old days—three years ago—most enterprise data networks were run on frame relay or ATM. But that's changing rapidly, as large users migrate to Multi-Protocol Label Switching (MPLS). Indeed, the carriers have made no secret of their desire to migrate users to MPLS over the next few years, whether or not the users want to go.

MPLS has several advantages. For example, you don't need to buy point-to-point permanent virtual circuits (PVCs) to connect ports through the carrier's cloud, as you do with frame relay. The platform offers any-to-any connections, greatly reducing configuration woes.

But its other key benefit—and the reason it is taking over the enterprise transport world—is the ability it gives users to prioritize different types of traffic by recognizing class of service (CoS) markings from the enterprise and allowing the segmentation of traffic into different virtual private networks (VPNs) within the same carrier cloud based on business unit.

Defining Class Of Service

CoS is a way of managing traffic in a network by grouping similar types of traffic together and treating each type as a class with its own level of service priority so that, for example, voice and transactional application traffic like enterprise resource planning (ERP) or customer service records (i.e., CRM) applications can take priority over email and large file transfers. CoS can also help in the utilization of a given fixed-port bandwidth, as when an enterprise assigns non-business Internet surfing to the lowest CoS.

The essence of CoS is prioritizing packets so that some get the fast lane while others get through whenever they can. Voice and video bytes have to get through quickly and consistently or the service is not usable. Data can withstand jitter and latency much better. Email can take all day (OK, make that 2,000 milliseconds).

Unless you have unlimited bandwidth, and no one really does, CoS is crucial to telecommunications convergence—you can't send everything over one pipe unless the network can recognize the stuff that has to get there right away and move it to the head of the line.

Vendor Comparison Travails

But for every yin there is a yang, and the yang of MPLS is that CoS capabilities and offerings vary among carriers in ways that make analysis and comparisons challenging. Some service providers offer six or more classes of service while some offer only three, and they all use different names (try comparing CoS 1, CoS 2, CoS 3, CoS 4, CoS 5 and CoS 6 to Platinum, Gold, Silver, and Bronze). And if that weren't bad enough, there is no commonality in the way the carriers price and charge for CoS.

AT&T, for example, offers multiple MPLS products: e.g. EVPN, AVPN, PNT and IP Enabled Frame Relay. They are all priced differently (in some cases even when they are provided using the

Start by modeling the "Day One" configuration of your network

same underlying infrastructure) and have different service level agreements (SLAs), and for most of its MPLS products, AT&T's pricing mechanism levies separate charges for required Class of Service functionality.

Customers of AT&T's managed MPLS product (EVPN) have to decide what Committed Data Rate (CDR) they require by Class of Service. For example, a customer with a T1 port who wants to use MPLS service for both VOIP and non-mission critical data applications might choose a Class of Service 1 CDR of 384 kbps (for use with VOIP), and a Class of Service 3 CDR of 384 kbps (for use with data traffic). The Class of Service 3 data traffic will be able to burst to port speed, but the Class of Service 1 VOIP traffic is not allowed to burst.

IP Enabled Frame Relay (IPeFR) is one of AT&T's un-managed MPLS products. IPeFR customers choose what Enterprise Permanent Virtual Circuit (EPVC) speed they require into the cloud. Once the EPVC speed (equivalent to the total CDR across all Classes of Service) has been chosen, the customer chooses the allocation of the EPVC across AT&T's four available Classes of Service from a standard CoS allocation list. The choice of allocation has no bearing on the charge, which varies only with the speed of the EPVC.

Other Pricing Models

The AT&T pricing models described above are not the only MPLS pricing models in the market. Some MPLS suppliers price their services like AT&T's EVPN—the customer pays for the specific bandwidth required for each CoS. A few do not vary pricing based on CoS requirements; the customer purchases an MPLS port, and then has access to all the supplier's classes of service, including real-time voice class of service. Yet other suppliers include the basic CoS (typically a best-effort CoS or equivalent) in the charge for the MPLS port, but charge customers extra for however much bandwidth they want allocated to higher level classes of service.

As opposed to levying additional charges by CoS bandwidth, some suppliers price MPLS ports depending on what CoS functionality is required. The lowest-priced port may provide only best effort CoS functionality, meaning that all of a customer's traffic would be allocated to a single CoS and all application traffic would have the same priority. Customers can buy more expensive ports that include a mix of higher-level classes of service, with the cost depending on the required mix (e.g., a higher cost may apply if real-time CoS is required to support VOIP). Under such models, the cost of the port may not be directly related to the bandwidth required for each CoS, but the price may vary depending on the mix of classes of services.

Finally, some combine more than one of the approaches described above, levying additional charges for CoS bandwidth only if a customer requires real-time CoS capabilities. For example, the customer may pay a higher port charge in order to have multiple classes of service (as opposed to allocating all traffic to a single CoS) and pay an additional charge on top of that, depending on how much premium, real-time CoS is required.

There are ways to justify additional charges for the premium CoS used for VOIP or other real-time application functionality. In order to provide the quality of service necessary for VOIP, the supplier has to manage the required voice CoS bandwidth more closely than non-real-time CoS, effectively guaranteeing that the bandwidth will always be available without other classes interfering due to bursting. True, not all carriers prioritize in the cloud—some simply over-provision capacity in the cloud so no congestion and contention occurs.

The major benefits of CoS prioritization are at the egress provider edge routers, where contention of port bandwidth is likely to occur. That's the only way to ensure that no congestion or performance issues restrict the traffic marked as real time COS, but such guarantees can only be made through robust engineering of the network and by properly servicing CoS queues, particularly at the provider edge router.

Conclusion

So how does a user compare bids in a market where the products are practically designed not to be comparable? Basically, you create a single model of how your network is most likely to be configured on "Day 1" and then price it using a universal pricing template, to ensure the bids you've received have apples-to-apples pricing.

If you want something more sophisticated, you can tweak the model to illustrate and price out different growth scenarios. It's not easy—actually, it's hellishly complicated, particularly for large, multinational networks. But it beats the alternatives—blind faith or sole sourcing□

Companies Mentioned In This Article
AT&T (www.att.com)

Hank Levine and Jim Blaszak are partners in Levine, Blaszak, Block & Boothby, LLP, (LB3) a Washington DC law firm that specializes in the representation of enterprise users in connection with their procurement of network-related services; before the FCC and other telecommunications regulatory bodies; and in disputes with service providers. LB3 and its consulting affiliate, TechCaliber, represent scores of large users, including about half of the Fortune 100. Jim and Hank speak and write widely on telecom sourcing and regulatory issues. See www.lb3law.com and www.techcaliber.com.