

# QOS over MPLS – the Complete Story

John Bartlett and Rebecca Wetzell  
BCR Volume 36, Number 2  
February 2006

As service providers entice growing numbers of enterprises to MPLS-based service offerings, customers need to know that MPLS alone cannot provide a complete, end-to-end quality of service (QOS). The complete QOS “story” has four “chapters,” covering traffic classification, traffic prioritization, bandwidth management, and traffic monitoring. An implementation that excludes any of these cannot provide optimal performance - and because no single vendor can deliver all four chapters, complete QOS-over-MPLS solutions require multiple vendors.

NetForecast has researched enterprise QOS-over-MPLS solution requirements and recently interviewed several enterprise users to learn how they have been faring in implementing solutions. We offer the following insight and guidance to help enterprises write their own successful QOS-over-MPLS stories.

## Why Are Enterprises Drawn to MPLS?

Everyone we interviewed is using MPLS-based IP VPN services, and they all told us they were drawn to MPLS because its mesh architecture solves traffic, IT management, and bandwidth management problems inherent in other technologies. For example, MPLS eliminates the complex PVC management tasks that plague frame relay, makes it easier to manage access link bandwidth, and enables more flexible bandwidth provisioning. Simplified management translates into fewer errors, less labor, and higher reliability.

The director of network analysis with a multi-billion dollar, global pharmaceutical firm told us he chose MPLS to ease the management burden of a major merger. “Because of the complexity of the move to the new [merged] environment, we had to touch almost every site,” he said, “and MPLS is an easier environment to manage than [frame relay] PVCs in terms of set speeds, bursting, and CIRs.”

QOS requirements for critical applications and the need to enable VOIP convergence also factored into MPLS migration for most of the companies we interviewed. For example, a network engineer for a large multinational software vendor told us that his

enterprise chose MPLS so that voice and video could be added to the network without sacrificing the performance of critical applications. “We wanted the ability to add voice and video with guarantees,” he added, “and [we wanted] also to outsource the management function.”

Either before or after their move to MPLS, each of the network planners and managers we talked with realized that they must implement all four of the QOS aspects that we identified above (see **sidebar: Review of Existing Solutions**)

## Chapter One: Traffic Classification

Step one for any QOS solution is to classify application traffic so it can be treated differently depending on its relative importance. Traffic is classified by identifying and marking it at the application endpoints or at intermediate devices such as switch/routers or dedicated QOS appliances. Marking consists of setting priority bits in either the Ethernet header (IEEE 802.1p) or in the IP header (DiffServ) of each packet, thus identifying the priority that packet should receive as it traverses the network.

Endpoints like PBXes or video conferencing bridges can successfully mark traffic to support enterprise business objectives, but users at their desktops may be tempted to “game the system” by classifying personal applications (such as multiplayer online games) to receive higher priority. Because network managers have to protect the network from such users, they often ignore endpoint classification markings.

Switch/routers can identify traffic based on known servers and/or known TCP or UDP ports, but many older switch/routers cannot identify traffic based on higher level protocol differences, such as discriminating among multiple HTTP applications that are traversing port 80, or identifying real-time traffic flows that are using dynamically allocated ports. Most newer routers have classification and marking features like Cisco’s Network Based Application Recognition (NBAR), although these tools can be hard to manage and can consume

valuable processor capacity at the expense of packet forwarding.

Dedicated QOS appliances, in contrast, are specifically designed to look deep into packets to accurately identify protocols as well as individual applications using those protocols, and to efficiently classify and mark WAN-bound traffic. The appliances must be placed at each network site, where they operate on the customer side of the WAN access router.

QOS appliances can enable more classes than MPLS services typically provide. Not only can they add finer control to multiple MPLS classes, they can also add control if the customer buys only a single class of MPLS service. In both these cases, the appliance classifies and prioritizes all traffic before it hits the WAN access router or the carrier's edge router, both of which simply forward the traffic.

This works well as long as the MPLS service vendor provides 100 percent of the committed bandwidth at all times, so marked traffic gets through in a timely manner and without interference from other traffic on the service provider's network. In these circumstances, bandwidth commitments should be verified with the provider when negotiating a service level agreement (SLA).

## **Chapter Two: Traffic Prioritization**

Once traffic is classified, we need to ensure that high priority traffic receives better service than lower priority traffic. An airport check-in counter provides a good traffic prioritization analogy where "high priority" first class passengers are checked-in faster than their coach class counterparts.

Managing traffic priority is most important where network speeds drop from high to low speeds, such as at the enterprise edge or at the edge of the MPLS service provider's cloud. However, implementing priority throughout the network (both within the enterprise and across the service provider's WAN) provides the most consistent high quality transport for critical applications.

Traffic priority, or class of service (COS) assignment, is determined by taking into account an application's importance to the business and its sensitivity to network-induced phenomena like latency, packet loss, and/or jitter. COS can be expressed in the packet headers in layer 3 (with

DiffServ code points, or DSCP) or layer 2 (with the IEEE 802.1p header bits), and most switch/routers support both mechanisms.

Traffic that is properly marked at the classification stage, by the endpoint, switch/router or QOS appliance, should have the correct priority across the entire network path. Care must be taken to intelligently map the COS assignments to the available MPLS service levels. MPLS provides several priority levels -- often referred to as gold, silver and bronze -- each with different service level agreements (SLAs) specifying minimum thresholds or service goals for latency, packet loss, and perhaps jitter. The bandwidth used by each traffic class must not exceed the negotiated bandwidth for that class in the MPLS WAN.

Three or four classes can adequately support the needs of performance-sensitive real-time traffic, and less finicky data traffic. Voice usually has highest priority, followed by video, and then data. Data traffic is often split into two classes, with higher priority given to interactive applications like Citrix or mission-critical applications like order entry or online sales -- while email, human resource applications or database synchronization receive lower priority.

Enterprises with many critical data applications and low WAN bandwidth connections, as well as those that simply want more control over application performance, will want a more granular QOS structure than MPLS alone can deliver. For example, the senior networking engineer with a multinational insurance company told us that, until he installed QOS appliances at each site, he felt constrained by his MPLS service provider's limited traffic marking. "We only had three classes, and we had no direct control over which applications got what priority," he said. "We had to be in constant contact with our provider to see the effect on the applications. Having the provider be responsible for QOS made it difficult to support."

Again, dedicated QOS appliances can help. Traffic managing appliances like those from Expand, Packateer, Peribit (now owned by Juniper) and others add levels of QOS granularity. In an MPLS mesh environment, an appliance is installed in front of each WAN connection. Each device both classifies traffic (if necessary), and then prioritizes traffic headed for the WAN to ensure each traffic class is treated appropriately.

### Chapter Three: Bandwidth Management

Bandwidth management is important to the QOS-over-MPLS story because applications can perform poorly or fail if too much bandwidth is allocated to or demanded by any class of traffic. Successful bandwidth management maintains an appropriate mix of high and low priority traffic (as explained in *BCR*, October 2004, pp. 16-22). Bandwidth management has two important aspects:

- The network aspect: It is essential to allocate enough bandwidth for lower-priority traffic to pass, even when higher priority traffic is using 100 percent of its allotted bandwidth. Traffic also must be “policed,” meaning it is denied when it would exceed the bandwidth allocation of its class, so that traffic in other classes can pass.
- The application aspect: The bandwidth needs and behavior of applications in the face of limited bandwidth must be understood and managed to avoid application disruptions and user frustration. We must not only size application-specific bandwidth appropriately, but also manage applications so they behave well given bandwidth constraints. For example, if we allocate 300 kbps for voice calls and that bandwidth allotment is completely filled, what happens if someone else tries to make a call? If we allow this call into the class where it belongs, it will degrade the quality of all the traffic in that class-- a clearly unacceptable outcome. To avoid it, the application can deliver a busy signal, or reroute the call via the PSTN. Although data applications generally degrade more gracefully than real-time video and voice applications, choices still must be made about how to manage data traffic that exceeds allocated bandwidth.

Bandwidth can be managed at the application source and within the routers. At the source, the number of simultaneous users can be limited in several ways:

- VOIP call managers and video gate keepers can be configured to deny calls that would exceed allocated bandwidth, or to route them to an alternate connection.
- Some applications have a license manager that can support a limited number of concurrent users. For example, 200 users may be licensed,

but only 100 of them can use the application simultaneously.

- In branch offices, limiting the number of end users that are allowed WAN access is a straightforward way to control use of fixed bandwidth applications like voice and video, although it doesn't help with bursty data applications. Data bandwidth demands also can vary significantly based on the task, the data being viewed and the user's skill.

Bandwidth can be managed within the switch/routers by setting each COS queue to forward traffic only up to a maximum data rate. When this rate is exceeded, all traffic in that class experiences delays. Although this protects the carrier's network core from excessive demand, it is not appropriate for real-time voice and video applications, and it may not be granular enough for some enterprises.

Once again a QOS appliance at the LAN/WAN boundary (assuming it also supports bandwidth management and traffic shaping) can alleviate these problems by differentiating among the applications within MPLS service classes. Traffic within a class or even a sub-class can be selectively “shaped” or throttled back to remain within pre-defined bandwidth limits. To enforce this granularity, each appliance must either pre-allocate bandwidth to each other site, or coordinate with the appliance at each site to negotiate additional bandwidth per class when needed, and ensure that incoming links are not overcommitted.

This granular shaping capability is important to the senior networking engineer of the global insurance company we interviewed. “Trying to fit all of our traffic into three priorities wasn't enough,” he said. “With [bandwidth management] you can get very granular in defining quality of service. In addition, you can define it not just by protocol or application type, you can define it down to which host talking to another host gets what kind of service. You have ultimate control over the quality of service that you provide, and that's a good thing.”

#### Managing Bandwidth by Reducing Bandwidth Demand

While traffic shaping can squeeze traffic into the available allocated bandwidth, caching and compression reduce the amount of bandwidth required. This is especially useful in parts of the world where bandwidth is still scarce and/or expensive. A number of the enterprises we

interviewed expressed concern about minimizing bandwidth use to control costs. One network manager we interviewed told us his company could buy 4 Mbps between the US and Mexico for \$40,000 a year, while 256 kbps to Uruguay cost \$140,000 a year.

Keep in mind that QOS only works correctly if the traffic demand is within the available bandwidth. If the traffic demand in a given class exceeds available bandwidth, all applications within that class suffer.

Caching and compression help reduce bandwidth demand without interfering with traffic classification. Caching devices hold data that is requested by a user at a remote location so subsequent requests for the same data don't consume WAN bandwidth. Compression reduces the size of files so they consume less WAN bandwidth.

Compression (and decompression) can be accomplished by cards and blades in WAN access routers, or by dedicated appliances. Most compression appliances either do classification themselves (in which case they look at the data before it is compressed), or they compress classified traffic and re-mark it with the same marking it had upon arrival. The edge and MPLS routers then look at the packet marking as usual.

#### **Chapter Four: Traffic Monitoring**

It's a fact of life that IP networks are in a constant state of flux, buffeted by continuous upgrades, additions and changes. This makes maintaining the right traffic loads, COS assignments, and bandwidth allocations for each class a dynamic exercise, like a high-wire balancing act in a hurricane. As if this weren't challenging enough, MPLS meshes enable any-to-any traffic patterns that make some traffic invisible to centralized network monitoring devices.

You can't fix what you can't see -- so if you don't monitor the traffic, you don't know what performance your network is delivering to which application types, and you can't optimize performance. This is the perennial argument for network management -- and it certainly applies to MPLS networks.

In legacy frame relay networks, PVCs behave like dedicated lines and are thus easily managed. The trade-off is a lack of flexibility: It is hard to shift

bandwidth and hard to manage all the PVCs. MPLS mesh architectures offer more flexibility, but the mesh also makes traffic monitoring more complex and requires additional tools.

You need to see the applications traversing the WAN, their COS markings, and how much bandwidth they are consuming -- and you need to track these metrics over time in increments ranging from minutes to years. Also real-time traffic requires real-time testing.

Why is this so important? Because users may complain -- but then again, they may not. If they are employees, they may not know performance could be better, and when an application is slow, they may either trudge along, or give up and go for coffee, hurting productivity. Customers or business partners, on the other hand, may simply abandon your site or service -- with or without complaining. In any case, it's important to get and stay ahead of user complaints.

Unfortunately, application servers don't provide the traffic monitoring information you need because it isn't their job to know where application traffic goes from the server, nor to understand network topology. In the network, routers can collect some of this information (e.g. NetFlow in Cisco routers), and send it to offline storage for subsequent analysis. But routers don't collect very detailed statistics, and router cycles are better used for routing.

Some MPLS services include access to a portal that displays bandwidth usage by class over time, but this information is limited to the service classes the provider offers. If you are using a traffic classifier and/or a traffic shaper, the service provider can't show you the finer granularity applied by these devices, even though the devices themselves can. Many appliances also can export performance data to a common data collection server for further reporting and analysis. For real-time traffic, some measurement tool vendors monitor VoIP and/or video traffic for quality, or create synthetic streams across the network to test real-time support (also see *BCR*, February 2005, pp. 18-23.).

"True" QOS-over-MPLS requires specialized traffic monitoring and reporting equipment or services that provide and integrate information at the right level of detail for all meshed links. The bandwidth being used at each traffic priority level needs to be tracked over time. Details about which applications are using the most bandwidth within a

class help IT managers determine how best to allocate available bandwidth and justify bandwidth increases when appropriate.

Monitoring and reporting have proven essential to the senior network engineer at the global insurance company we interviewed. “We have new applications coming on all the time,” he said. “We have a development group here that builds their own applications, and they don’t tell us in advance, ‘we’re putting a new application out there, and here’s what you have to look for.’”

“So when they deploy these applications globally,” he continued, “all of a sudden we get calls saying the network is slow, and we have to hunt down the problem, and find out it’s a new application that’s not performing well. Once we find out what the application is, and how best to classify it, we can go ahead and make the change and resolve the problem pretty quickly.”

If you are reluctant to embrace the need for ongoing detailed monitoring, be aware that you will need to buy more bandwidth than you might otherwise need. With huge pipes and no congestion, none of this monitoring and management is required. But if you have to pay \$140,000/month for 256Kbps to places like Uruguay, you will probably want to manage every last bit.

### **Benefits of a Complete QOS-over-MPLS Story**

The director of network awareness and analysis at the multinational pharmaceutical company we interviewed found that by implementing a complete QOS-over-MPLS solution he realized more than 10 percent savings in network services and operations staff resources, and he was able to maximize bandwidth investments by “squeezing more out of every connection.” He also discovered that “bandwidth cost went down dramatically, but our budget stayed the same, so we were able to buy more bandwidth for the same amount of money.” In addition, knowing what traffic is traveling over the network made the network easier to manage and enabled network convergence (voice & video combined on the data network).

According to the network engineer, IT Operations at the multi-national software vendor, “the primary benefit of [MPLS] has been the traffic classification and prioritization, and the ability to set policies for different applications,” he said. “Another huge benefit was the implementation of

WAN compression, which saved us lots of bandwidth and reduced our cost.” QOS-over-MPLS also guaranteed performance for critical applications, and enabled high quality voice and video, he added.

The primary benefit the global insurance firm experienced after implementing QOS-over-MPLS was guaranteed performance for critical applications that resulted from being able to clearly define and classify service levels. They also found increased network visibility a big benefit, along with the ability to troubleshoot problems. According to their senior network engineer, “It’s faster and easier to troubleshoot problems because we have insight into the data flows in our network.

### **Conclusions**

Achieving QOS-over-MPLS entails much more than simply signing up for MPLS service: It is a challenging assignment that requires understanding the whole QOS story without skipping any chapters. Although achievable, “true” end-to-end QOS-over-MPLS demands careful integration and tuning of components from multiple vendors.

### **Sidebar: Review of Existing Solutions**

Each of the four functions in **Table 1 (QOS-over-MPLS Solutions)** is required to insure a complete QOS solution across the network. Routers and MPLS service providers each offer portions of the solution, but none are complete. Instead, each needs to be augmented with additional monitoring and/or QOS appliances to fill out all the requirements.

Note that router vendors and MPLS service providers each will claim to provide bandwidth management. What they each mean is that the bandwidth for a class of traffic will be limited to a predetermined value. This feature ensures that the network can continue to support high priority traffic with the right characteristics (low packet loss, low latency and low jitter). But it does not mean that applications will be well served. Traffic shaping can help this problem by smoothing out demand to meet the available bandwidth.

At the highest level, planning is necessary to ensure that sufficient bandwidth has been allocated to meet the user and business needs of the enterprise. After the network has been set up,

ongoing monitoring is necessary to keep it tuned to perform properly.

**Table 1. QOS-over-MPLS Solutions**

	<b>Traffic Classification</b>	<b>Traffic Prioritization</b>	<b>Bandwidth Management</b>	<b>Traffic Monitoring</b>
<b>Router</b>	Limited to Layer 3	Yes	No	Limited
<b>MPLS Service Provider</b>	No	Yes, only 3 to 5 classes	No	Some provide monitoring, limited to classes they offer
<b>Monitoring Appliances</b>	No	No	No	Yes, excellent granularity and detail
<b>QOS and Traffic-shaping Appliances</b>	Yes	Yes	Yes, through traffic shaping	Yes, through distributed monitoring and central control