

Thinking About

Making the Transition to MPLS

Why Consider Multiprotocol Label Switching (MPLS)?

Many organizations are considering a move from Frame Relay and ATM to Multiprotocol Label Switching (MPLS)-based network services. MPLS VPNs provide the privacy and security of a Frame Relay or ATM network, yet they allow the inherent any-to-any connectivity and routing flexibility typical of an IP network. It's the best of both worlds.

In addition, MPLS VPNs provide three other important elements:

1. Quality of Service (QoS) levels that are critical to application convergence and they also support voice, video and data on one network.
2. An MPLS VPN is a natural platform for dynamic disaster recovery. Based on IP routing, it's simple to support multiple data centers in a load sharing or primary/backup scenario.
3. MPLS VPNs simplify the data center architecture. Unlike with Frame Relay and ATM, you are not constrained by managing many point-to-point direct connections to remote sites.

MPLS Layer 3 VPNs are based on RFC2547bis. This standard defines mechanisms to keep individual VPNs completely private and separate from each other. While the data traverses a shared infrastructure, customers can be assured that their VPNs are maintained and monitored as separate entities, much like a Frame Relay or ATM network.

The following summarizes the benefits of an MPLS VPN:

- Any-to-Any IP Connectivity – Optimal Routing without PVCs
- Improved Latency – Avoid tandem routing through a hub (offload hub router)
- Any IP Address Scheme – Intranets and extranets
- Circuit Consolidation – Eliminate aggregation layer, if Layer 3 VPN, for all sites
- Cost-Effective – More cost-effective than Layer 2 services especially as connectivity requirements increase

- Diversity Via IP Routing – Simplified Disaster Recovery and inherent redundancy
- Ease of Network Expansion
- High Speed Access Support – Up to OC192
- Access Technology Agnostic – Frame Relay, ATM, PPP over DS0-OC48; Ethernet, DSL, etc.
- IP Class of Service
- Provider-Based IP VPN – No requirement for CPE-based tunneling and encryption equipment/software/ overhead nor PKI management
- Easy to Support Plug-ins – Video and Voice Gateways, Remote Access, Network-Based Firewall

Migration Considerations to Think About

- Complexity
- Critical applications performance
- Traffic management
- Security in peer-to-peer model
- Design of routing architecture

Migration Considerations to Think About

The advantages of MPLS are compelling. In the course of migrating to MPLS, organizations should consider the impact of both the network migration itself and the ease of operating and managing in the new environment.

AT&T research shows that among its customers who are considering migration from Frame Relay or ATM, the primary consideration is increased complexity. This complexity issue itself then breaks down into four primary drivers, each of vital importance to keeping business operations running smoothly.



The four drivers include:

- Critical applications performance
- Traffic management
- Security in a peer-to-peer model
- Design of the routing architecture

The remainder of this paper examines these topics and discusses how to address them in your migration plan.

Assuring That Your Critical Applications Perform Well

Designing networks for quality of service deployment is a high stakes exercise. It is particularly important for Voice over IP (VoIP) applications. In traditional data applications, the protocols and applications themselves can be pretty forgiving. Some of the features of a frame network, such as traffic management and protocol features like TCP sliding windows, are simpler on an MPLS network because the traffic management headaches are moved into the network.

In more sensitive applications such as VoIP and IP Telephony, there is often no ability to tolerate oversubscription, and the voice traffic volumes (call volumes) can be fairly unpredictable. This puts additional burdens on network planning personnel.

“It’s a new world out there with voice traffic,” says Tom Siracusa, Director AT&T Labs, VPN Strategy. “You have to be pretty precise in your analysis of the traffic you have today, and fairly granular in your plans. For example, you need to understand the number of concurrent voice calls expected over a given WAN link. It’s not like TCP/IP and bulk file transfers, where sliding windows and retransmissions can take care of things in the background.”

The first step to assure your applications perform exactly as your business requires is a thorough traffic analysis, resulting in a baseline of your current traffic patterns. Fortunately, a range of hardware and software tools are available to help with that traffic analysis. One option is to instrument your current network and build your applications’ profiles from there. Another option is to bring in professional services to put equipment on your network on a temporary basis to gather a snapshot that can be used for detailed planning. The other side of the planning process is to take the traffic information (by application or application class) and then build the MPLS traffic profiles, based on the number of classes of service you have available. This is more complicated than the PVC designs from Frame Relay in that you may be considering full mesh designs with real-time traffic, such as voice, that is not very tolerant.

Whichever option you select, the critical output is an accurate profile of your existing traffic and an accurate model of your new applications. These profiles are the fundamental data set for constructing the Class of Service (CoS) regimen in your new network.

Improving Traffic Management for Application Performance

A good network design is the foundation for good traffic management, built upon the solid traffic analysis discussed earlier in this paper.

Supported by this foundation, your current applications should run as well – or better – on your new MPLS network, and your new applications should perform exactly as designed.

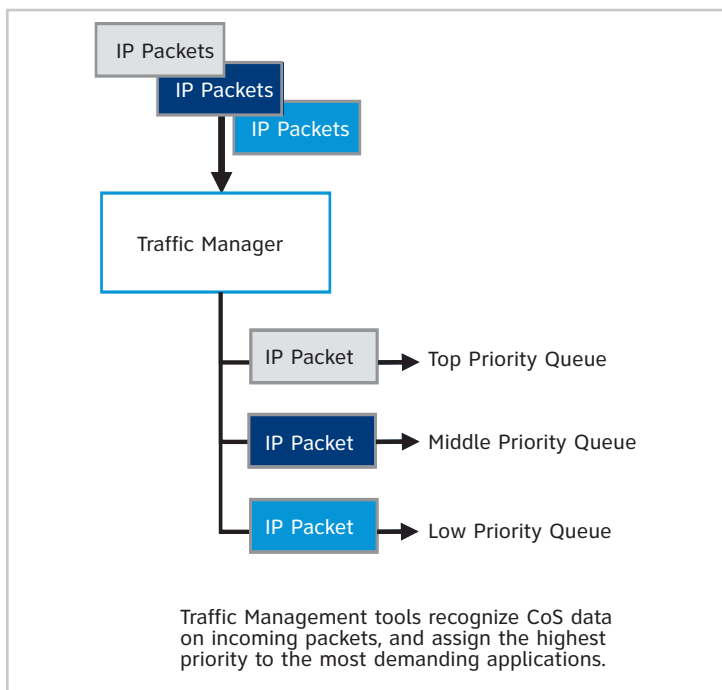
QoS techniques and network performance monitoring tools will ensure that you maintain consistency of application performance, and that you’ve got insight and reports to trouble-shoot problems quickly.

Setting parameters on the required performance metrics with associated SLAs generally requires that you purchase a specific CoS to serve the traffic. To separate out and define this traffic with a policy so that the MPLS VPN service can identify it, your network needs to be instrumented to achieve visibility across the entire domain, from the smallest branch to the largest data center. You can choose from the following tools:

- A system purchased as an appliance by the enterprise
- WAN-access router software with QoS capabilities bundled in
- Premises-based equipment purchased as part of a managed network service from the MPLS VPN service provider
- A network-based managed service offered by the MPLS VPN service provider

These tools enable you (or an outside expert that you hire) to first baseline your applications by monitoring the bandwidth each consumes, their behavior and requirements. Then, you can use the tools to classify traffic based on user group, application, metric requirements or other criteria and prioritize them through the network accordingly. You can automatically set policies to guarantee minimum levels of bandwidth for mission-critical, delay-sensitive traffic and limit bandwidth to other applications, such as certain peer-to-peer traffic.

Traffic Classification



Achieving Security in a Peer-to-Peer Model

MPLS VPNs tend to stir up security concerns, particularly given the openness of operating in a peer-to-peer model. These concerns can largely be put into two categories:

1. Is an MPLS VPN really private?
2. Do the meshed characteristics of an MPLS VPN necessitate that I change my security architecture?

The following topics will be discussed: MPLS VPN privacy, distributed network-based and endpoint security, traffic management and routing architecture design.

MPLS VPN Privacy

The MPLS VPN IETF standard defines mechanisms to keep each VPN completely private and separate from other VPNs. While the data traverses a shared infrastructure, enterprises can be assured that their network is maintained and monitored as a separate entity. Much like their Frame Relay and ATM network service counterparts, MPLS services behave as a segregated “tunnel” through a shared infrastructure.

Rather than provisioning PVCs and Data Link Connection Identifiers (DLCIs), your service provider will build virtual routing and forwarding (VRF) tables to create customer separation. A VRF is created by the carrier for each individual VPN. Each VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table. A VRF containing this set of data exists for each VPN site attached to each carrier’s Premises Edge (PE), keeping each VPN connected to each PE completely segregated.

Distributed Network-Based and Endpoint Security

Traditionally, most enterprises have kept very strong separation between their private intranets, their Internet access and their extranet connections to partners and suppliers. Over time, this separation is beginning to diminish.

This is because businesses need certain applications, such as VoIP, to seamlessly span all three. Meanwhile, mobile workers are using their work laptops for connections to public networks. In such a dynamic environment, centralized security starts to become obsolete. The inherent mesh of an MPLS network exacerbates this concern.

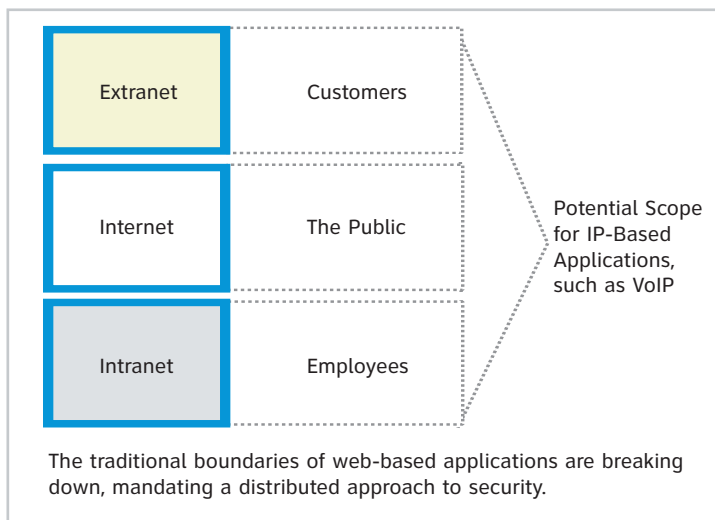
Traditional security models, though, tend to parallel the centralized, hub-and-spoke configuration of Frame Relay networks. Left as is, this could be problematic. For example, a security appliance at a centralized site filters traffic coming only through the “front door” of the enterprise. However, with direct site-to-site MPLS links, a virus/worm outbreak at a remote site could spread to many sites before the appliance were to discover it and start taking remedial action.

Therefore, shifting to a distributed security model is critical. Included in this shift will be the use of the network as one of the key security elements in the total environment, taking advantage of the network’s ability to monitor activity across a vast number of endpoints, enabling it to spot and correlate anomalies at the earliest stage of emergence.

Regardless of the type of VPN service you use – Frame Relay, Internet-based IP VPN, MPLS-based VPN – it is advisable to have both a network-based and an endpoint security architecture. Intrusion prevention systems can be installed directly on WAN access routers or in the form of standalone WAN-edge appliances to scan traffic for malicious signatures before granting access. Dynamic access to the latest signatures is available from router and appliance vendors or antivirus companies.

If you prefer to outsource the management of your endpoint security, scanning and remediation capabilities are also available in service form.

Security in an Enterprise MPLS Network



Reducing Complexity Through Routing Architecture Design

Another issue on the subject of complexity is routing. MPLS VPN services support either eBGP or static routes between the CE and PE.

Most business customers are familiar with interior gateway protocols, such as Open Shortest Path First (OSPF) and Cisco Enhanced Interior Router Gateway Protocol (EIRGP).

When using an MPLS-based service, however, your WAN access routers will be operating at the IP layer, peering at Layer 3 with the carrier’s MPLS edge routers.

The best protocol for peering with the carrier’s MPLS network is external Border Gateway Protocol (eBGP) to exchange route reachability information among different routing domains. While using BGP requires that the enterprise network manager learns a little about BGP, learning enough for using an MPLS service is not at all difficult.

In addition, your carrier should provide field experts who can offer basic advice on setting up your premises equipment. Alternatively, a managed service would shift all responsibility for BGP configuration into the hands of the provider.

The Strong Benefits of Using External Border Gateway Protocol (eBGP)

- Allows load balancing across many MPLS network topologies using BGP
- Supports several different redundancy options for carrying primary and backup routes
- Treats the Carrier Backbone with non-zero cost to optimize routing for all routes, including “backdoor” routes
- Supports Outbound Route Filtering (ORF) which allows a CE to dynamically signal a PE to only advertise certain route prefixes to it
- Offers several capabilities to provide enhanced and very specific routing control
- Allows route dampening option that aids in route flap protection; BGP does dynamic timer negotiation at BGP session initiation for the keep-alive and hold down timers
- Provides one of the key protocols (MP-BGP) used by all carriers (including AT&T) that supports RFC 2547-based MPLS networks
- Is the most efficient, scalable and reliable protocol for handling large numbers of routes and network state changes
- Uses a fairly simple decision process when deciding which path to use for a route, minimizing administrative complexity

Conclusion

Enterprises that moved from private line to Frame Relay networks over the last 10-15 years initially had many of the same concerns we hear associated with MPLS today. The primary motivator was cost savings, which they did realize. They also discovered that the application performance, the security and the network management issues were all manageable and indeed are now the gold-standard to which they are comparing this next migration.

The benefits of operating in a fully converged, meshed architecture will quickly overtake any of the perceived concerns and, undoubtedly, MPLS VPNs will be the de facto standard to which the next major technology innovation is compared.

“The MPLS network is paying off – we’re seeing enhanced employee productivity, easier administration, improved reliability, and measurable improvement in the total cost of ownership.”

Andy Daudelin, AT&T Vice President, IT Operations

Action Items

Solutions

Shifting from hub-and-spoke to mesh topology	Move to a distributed security model by embedding intrusion prevention scanning in distributed WAN access routers
Commingling traffic with other businesses on IP network	Keep traffic secure by circuit-like partitioning using label-switch paths
Maintaining performance SLAs	Buy premium service class for premium traffic; use traffic management tools to mark priority traffic and to monitor SLAs
Moving to BGP at the WAN edge	Learn new Border Gateway Protocol (ask carrier for help and/or consider managed network service)
Justifying costs	Utilize one network for voice and data less expensive to run, over time, than two; consider payback of new integrated multicast, VoIP, call center, and presence applications

For more information contact an AT&T Representative or visit www.att.com/business.

