BLUE COAT

Security Empowers Business

OPTIMIZING APPLICATION TRAFFIC ON MPLS WAN LINKS

SOLUTION BRIEF

MPLS (Multi-Protocol Label Switching) is a common type of service offering from global telecom providers that removes much of the cost and complexity of more traditional point-to-point networks. The change from point-to-point to any-to-any or cloud networking presents new application delivery challenges, requiring a response from services that depend on the packet delivery infrastructure.

Performance Challenges

Businesses have been switching to MPLS WAN links because they remove the cost, complexity and routing challenges associated with maintaining dozens or hundreds of point-to-point leased lines or IP VPNs over the Internet. Carriers offering MPLS benefit from economies of scale, peak load balancing between customers and higher service margins from outsourced network management. This allows carriers to compete efficiently with the do-it-yourself, direct to the Internet networking alternative. However, the any-to-any automatic networking offered by MPLS, while having distinct advantages to traditional connectivity, presents new challenges that ripple through higher-order services that depend on the packet delivery infrastructure. These include the intersection of other key IT priorities – such as server consolidation, storage networking and use of real-time multimedia – with the distributed nature of MPLS. To compensate for these challenges, carriers often provide differentiated service levels, either by application/port or between two endpoints. Yet this too presents its own management complexities. This is especially true in today's highly distributed enterprises where the Internet is an integral part of business operations. In these environments, four key issues emerge:

1. Lack of application visibility makes prioritization difficult. Internally-hosted business applications are increasingly becoming web-based, making them more easily accessible to a distributed workforce via a standard web browser. However, because MPLS is limited to layers 2 and 3 of the OSI model, everything coming through port 80 and 443 looks the same, making it difficult or impossible to prioritize web applications using the built-in MPLS service levels.
2. Server consolidation creates MPLS bottlenecks. Server consolidation, by bringing servers out of branch offices and into regional datacenters, creates islands of dense network activity in the MPLS cloud. While endpoint offices benefit from distributed access to applications over higher bandwidth with MPLS, the backup and synchronization traffic between datacenters becomes more concentrated but over smaller links. The result is either very expensive datacenter to datacenter exceptions to MPLS policy, or a return to more expensive leased lines between datacenters.
3. Encryption is required on shared networks like MPLS. Unlike leased lines, MPLS networks are not strictly private; to preserve routing and QoS, carriers inspect network flows and co-mingle data between organizations. Because the privacy and security guarantees vary between carriers, especially internationally, where carriers must often sub-contract, regulatory compliance often requires blanket encryption of sensitive data, often with SSL. Yet encryption obfuscates traffic, making it difficult to prioritize, and accelerate.
4. Backhauled Internet traffic is a growing problem. For security reasons, most organizations centralize Internet access and back-haul Internet traffic across WAN links to remote users. As more and more business-critical applications are hosted on the Internet and the number of business-appropriate websites increases, the amount of MPLS traffic dedicated to backhaul becomes a significant cost. Further, the added distance to the gateway also adds latency, decreasing the performance of Internet-based applications.

BLUE COAT

Security Empowers Business

There are options available to address these challenges. One option is to make adjustments to the MPLS service levels. However, that is limited only to the traffic visible and understandable to network routers. Neither the carrier nor the internally managed router can distinguish between critical and trivial web traffic, or open tunneled backups or encrypted web applications. Indeed, the MPLS service-level architecture already addresses most of what can be accomplished with Layers 2 and 3 of the OSI model. To further accelerate applications, align network priorities with organizational objectives, and provide agility to application delivery, a solution higher in the network stack is required to overcome the challenges of MPLS networks and position the enterprise to exploit its unique advantages.

Solution Strategy

Any solution that improves and extends MPLS deployments must act above the packet delivery infrastructure to address the application-level complications of cloud networking. Network proxies are one such solution: application, user, and network aware, they provide the logic to optimize the underlying network for the applications running on it. Secure WAN optimization proxy appliances at each end of the corporate WAN link provide complementary performance, security, and control capabilities. The requirements of such a solution include:

1. Policy-based logic for the network layer. WAN optimization appliances, deployed as a higher level overlay architecture, must be able to provide granular visibility into user-application sessions – including the ability to distinguish between recreational Internet traffic, web-based business applications backup traffic, and time-sensitive video or VoIP. With this visibility, the appliances should apply policy-based QoS controls to prepare the traffic for optimal class assignment on the MPLS link.
2. Acceleration to the last mile. Just as important, but often overlooked, these WAN appliances must help optimize the performance of application traffic that now may travel much farther through a cloud than in the point-to-point legacy network. This logical and physical network distance increases latency that ruins application performance. Worse, congestion and latency on the "last-mile" links connecting branch offices to the MPLS cloud further exacerbate latency, and can render MPLS investments ineffective.

By applying a holistic view to WAN performance, application performance policies can be assigned end-to end, at every point along the access WAN and MPLS tunnel to ensure optimal application response time for all remote users. By doing so, IT untangles the decision to use MPLS from other key initiatives. The capabilities required by WAN optimization appliances to address these performance challenges are summarized below.

Bandwidth Management

- A WAN optimization appliance can assign bandwidth and prioritization markings to traffic at the enterprise edge, based on specific characteristics (application, transaction type, user, etc.), effectively providing application-layer intelligence to the packet delivery network. This type of granular policy control can help companies recover wasted MPLS bandwidth and apply it to the users and applications that matter most.

Protocol Optimization

- Although MPLS often means improved bandwidth to the endpoints, the migration to a cloud network often means longer logical and physical connections than dedicated point-to-point leased lines. Moreover, the improved connectivity between sites encourages the use of applications that were not designed for WANs, such as CIFS for file sharing, MAPI for email, and others. Optimization appliances can intercept traffic and modify application protocol behavior to overcome the effects of latency with complete transparency to the server, end user and network management tools.

Application-level Security

- In order for bandwidth management and protocol optimization to be effective, the appliance must be able to understand the user-application interaction natively. That means comprehensive integration into existing authentication and audit infrastructures, as well as understanding application protocols and how they work across a variety of different network conditions.

BLUE COAT

Security Empowers Business

SOLUTION BRIEF

SSL Control

- SSL encryption is a requirement on semi-public networks like MPLS, but performance and control should not be sacrificed for privacy. An application delivery appliance must be able to provide visibility into all SSL sessions that are legally and operationally appropriate to intercept. By doing so, it can assign proper service levels to SSL traffic (including no service) and expose the traffic to protocol optimizations and caching – and then re-encrypt it for privacy.

Object Caching

- By caching user files, web content, and video closer to the user, optimization appliances can dramatically reduce the bandwidth required to provide high-quality services over MPLS. Transparent to the user and administrator, object caching is the most effective optimization technology, reducing user wait times and bandwidth to virtually zero for cached content.

Byte Caching

- Byte caching augments object caching by storing common network patterns, enabling it to partially cache almost any file, even as they change. This library of cached network pieces acts like a customized compression algorithm for your network, siginificantly reducing both bandwidth and latency required to service remote MPLS users.

Compression

- Inline compression can reduce predictable patterns even on the first pass, making it an ideal complement to byte caching technology.

How Blue Coat Can Help

Each of the techniques described above can help optimize MPLS WAN value and performance in general, and the performance of business-critical applications in particular. A few of these technologies are available from many vendors. But only Blue Coat delivers all MPLS WAN optimization capabilities through its patented MACH5 acceleration technology in both an appliance and virtual appliance format. In addition, Blue Coat complements these WAN optimization features with industry-leading web security and granular policy-based controls; all in a completely transparent, secure and resilient application delivery framework.

Blue Coat Systems Inc.
www.bluecoat.com

Corporate Headquarters
Sunnyvale, CA
+1.408.220.2200

EMEA Headquarters
Hampshire, UK
+44.1252.554600

APAC Headquarters
Singapore
+65.6826.7000