MPLS VPNs: The Foundation for Emerging Apps

A recent survey of nearly 400 Webtorials subscribers revealed that the growing use of real-time and streaming applications is driving and sustaining enterprise MPLS VPN service deployments. Why? MPLS's quality of service (QoS) and inherent multipoint routing characteristics, combined with virtual private network (VPN) security, make MPLS VPNs a strong foundation for privacy-sensitive corporate application traffic with requirements for predictable, low-latency, high-bandwidth network performance.

For additional information on the Webtorials MPLS VPN survey methodology and respondent demographics, visit Webtorials.

In addition to traditional business data applications, the applications increasingly running over MPLS VPN services are collaborative in nature. MPLS VPN survey respondents reported high levels of voice over IP (76%), unified communications (51%) and streaming audio/video (43%) usage on their networks.

These applications all inherently require low-latency network conditions. MPLS's ability to prioritize delay-sensitive packets ahead of others and to support direct site-to-site connections for minimal delays, then, makes the network a particularly good fit. (see **Figure 1**).

Figure 1. Importance of MPLS VPN Network Characteristics

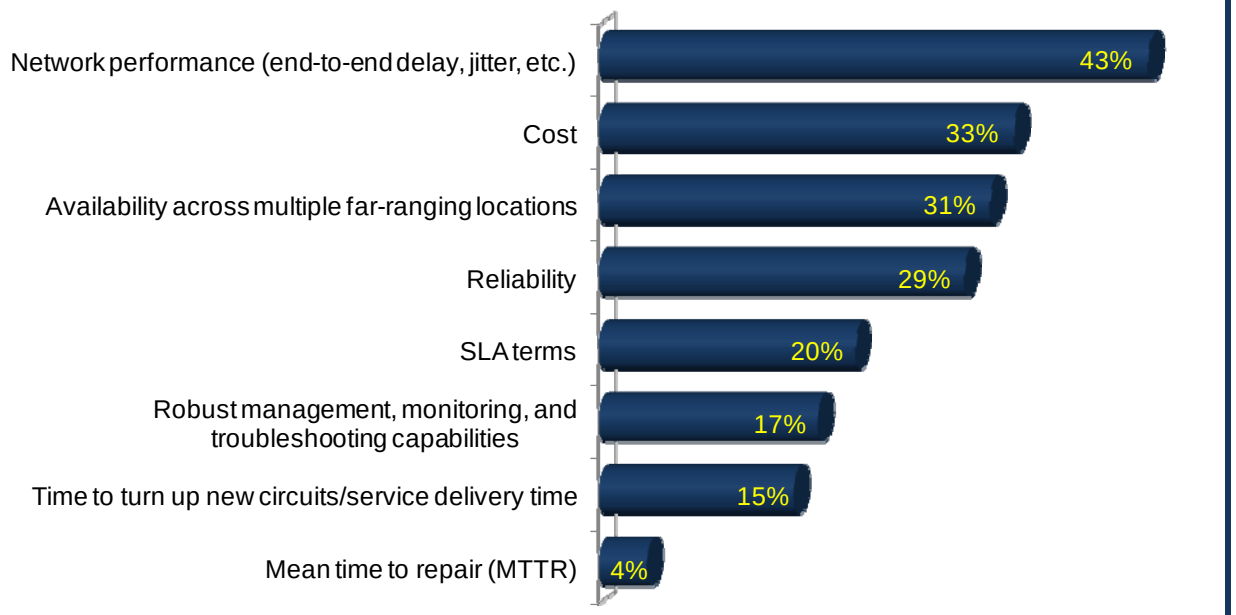

Which of the following do you consider the TWO MOST IMPORTANT MPLS VPN service attributes?

User Behavior

One reason for the uptake in collaborative application deployments is the general shift toward instant, dynamic communications and application sharing. The demand for these business functions has been fueled, in part, by consumer trends toward using mobile and social networks.

Because of mobility, consumers and now business users have come to expect the ability to communicate and access information anytime and anywhere. Similarly, though the roots of social networks like Facebook initially took hold in the consumer market, these networks are no longer reserved just for teenagers. For example, the U.S. government supports a Federal Chief Information Officers (CIO) Council, which uses social networking to exchange ideas and experiences in a community fashion. Instant collaboration capabilities are expected; in business and government environments, a level of privacy not offered by public Internet access connections is also necessary.

While the business climate is dominated by e-commerce and global expansion, it is also experiencing a weak economy and across-the-board budget cuts. Businesses require, now more than ever, that employees do more with less. That mandate, combined with the latest technological advances in broadband networking and performance management, has created a perfect storm for the use of electronic collaboration. When appropriately supported over the right kind of network service – one with low latency, direct inter-site connectivity capabilities and privacy – collaboration helps contain costs while improving response times and productivity.

Economic Impact on Application Deployment

Many of the latency-sensitive, collaborative applications described can also help corporate budgets by replacing more expensive business processes that involve travel. Much of this electronic activity will take the form of on-demand audio conferencing and desktop videoconferencing. Increasingly, too, department heads will communicate with their teams using streaming communications across wide geographic boundaries.

At the executive level, telepresence technology (described more fully in the next section) affords the ability to connect "in person" at a moment's notice. Businesses can meet with customers, partners, suppliers, coworkers and others to build relationships, strategies and plans economically and conveniently across international borders. The direct any-to-any connectivity capabilities of MPLS VPNs enables a given organization to add a partner or supplier to the routed network quickly, without having to pre-configure virtual circuits from an existing enterprise location(s) to the new location(s).

The latest collaborative tools have the potential to deliver bottom-line productivity improvements and overall lower expenses. It's inevitable, however, that the ever-tightening budgets of some companies will get in the way of some collaborative deployments, despite their potential to allow users to do much more with much less.

Unified Communications Outlook

Unified communications applications include a range of capabilities that initially have revolved around IP telephony. One has been the unification of different types of electronic mailboxes (voice, email and fax, for example). Such capabilities pare the time required for checking multiple mailboxes and potentially missing a high-priority message that could make or break a deal.

Combined with presence (location) and fixed-mobile convergence (FMC) applications, unified communications becomes a platform for instant, dynamic communications. One useful application is having a single phone number reach across multiple wired and mobile phones (a feature often called "single-number reach" or "twinning"). Single-number reach takes a big bite out of the time and productivity wasted playing the proverbial game of telephone tag.

From there, Web conferencing services, which allow users in various locations to view and manipulate common documents, are also taking off. Particularly when used with their voice over IP (VoIP) components, these apps require a highly tuned network foundation with minimum delay, jitter and packet loss and highly reliable and consistent throughput.

Telepresence, mentioned earlier, can be considered the mother of all collaborative applications. Although still in early deployment stages, telepresence made a respectable showing in the Webtorials MPLS VPN survey responses. About 22% of respondents said they were currently using or anticipate using telepresence over an MPLS VPN service during the next year.

Telepresence is a life-size conferencing application that uses 3D images and high-definition video to simulate in-person conferences, often among as many as 10 people on each side of the network connection. Telepresence technology, on the market for a few years, requires plenty of bandwidth (about 15Mbps), direct connectivity and high performance.

To date, telepresence has initially been a bit of a niche service for deep-pocketed companies, as equipment can cost $300,000 and up per site, depending on system used. However, two factors are likely to drive greater use of telepresence:

1) An economy seeing slashed travel budgets but still requiring face-to-face connections for building relationships and optimizing productivity. Desktop and room videoconferencing have filled the bill in some instances for serving this purpose; however, telepresence offers a whole new category of conferencing with its lifelike experience. This capability is particularly important for companies wishing to build very close working partnerships with customers or business allies at the executive level.

2) The growing availability of telepresence as a managed service from network service providers. In service form, telepresence offers a pay-as-you-go billing model, which is likely to be more financially palatable to capex-shy companies in a down economy.

Impact of New Apps

The rise of high-bandwidth, performance-sensitive collaborative applications is motivating enterprises to find ways to unclog the traditional last-mile capacity bottleneck – both with additional bandwidth and with traffic management capabilities. For example, Ethernet in the access network (a.k.a. Metro Ethernet) has become a common access service for connecting to the MPLS VPN service cloud. More than half (55%) of the Webtorials MPLS VPN survey respondents use Metro Ethernet, making it the most popular MPLS VPN access service. Metro Ethernet services are generally available in locations where fiber to the premises is available, and they can support several hundred megabit-per-second speeds.

In a unified IP environment, however, applications with different types of behavior and network requirements co-exist side by side. As organizations continue to grow increasingly distributed, whether domestically or internationally, increasing volumes of this unified traffic are traversing the WAN. As such, controls are needed to make sure each application type sharing the same "pipe" – even a big pipe – has the resources it requires to perform optimally. MPLS VPNs inherently provide some of these controls through the use of their packet-prioritization-based QoS.

Tuning each traffic flow within the unified traffic stream across the WAN is further driving the adoption of special products and services known by such monikers as *WAN optimization*, *application acceleration*, *traffic management* and *bandwidth management*. These capabilities use a technology called *deep packet inspection (DPI)* to determine the application type of each packet. Once identified, the packet can be treated in accordance with QoS policy.

For example, enterprises might wish to cap the rate at which certain types of traffic flow (or the percentage of overall available bandwidth they consume). This function is called rate limiting. Rate limiting is frequently used for known "bandwidth hogs" such as peer-to-peer (P2P) file sharing applications, which, if left unmanaged, can quickly fill up an entire link and degrade other applications. Enterprises can also configure the network to ensure a small, but consistent amount of bandwidth to session-oriented traffic such as voice and Citrix thin-client applications that can't endure packet loss.

WAN optimization capabilities can be acquired and managed by the enterprise in the form of customer premise equipment (CPE) or procured as a managed network service from a network service provider. Again, in a budget-constrained economy, the managed service option offers a pay-as-you-go alternative to capital layouts and associated training requirements.

In a down economy, the WAN optimization/bandwidth management capabilities described also play a distinct cost-savings role. They allow enterprises to postpone investments in additional WAN bandwidth by helping them squeeze as much utility out of their existing bandwidth as possible. Controlling traffic behavior on a flow-by-flow basis minimizes capacity waste while ensuring that each application, with its unique resource requirements, behaves optimally. MPLS VPNs inherently provide some of these controls through the use of their packet-prioritization-based QoS.

Security Service Trends

As mentioned earlier, enterprises and governments need the privacy afforded by MPLS's virtual private networking capabilities. MPLS VPNs create private customer tunnels in the form of virtual routing and forwarding (VRF) tables. VRF tables for each customer are built on each MPLS provider edge (PE) router at the boundary of the provider's MPLS network and the user's access network, and they partition one customer's traffic from that of others.

There are numerous layers of security, though, and even in a down economy, enterprises should not scrimp on security. The various security types pertain not only to access control and user authentication for protecting confidentiality of data. They are also needed to protect against malware and denial-of-service (DoS) attacks, which affect the availability of the network and, thus, the organization as a whole.

Most enterprises using MPLS VPN services also use an Internet gateway service so that internal MPLS users can also gain access to public Internet sites. So protection against picking up malware on the Internet and infecting the corporate network is necessary.

Survey respondents currently using MPLS VPNs most commonly used network-based firewall (62%) and antivirus (46%) services. (See **Figure 2**.) Firewall services can range from the enforcement of corporate outbound and/or inbound access policies to blocking access to known malicious sites where spyware lurks or against-the-law activities, such as gambling, take place.

Security applications, like WAN optimization capabilities, can exist as a service or be deployed and operated by the enterprise. Highly distributed enterprises will find it increasingly cumbersome to deploy and manage security CPE such as firewalls as they add more sites and begin to use software as a service (SaaS) applications hosted in the Internet, rather than in their own data centers. The uptake of SaaS and other application hosting should drive firewalls closer to where the data is stored, which is in the network. Similarly, antivirus services are growing in popularity because they serve to stop infections out in the public network, before they are anywhere near the enterprise's own private network infrastructure.

Intrusion detection services (IDSs) are catching up to firewalls and antivirus services faster than content filtering and anti-spam services. Of those using MPLS now, 42% are using network-based IDSs and 14% are planning to deploy them. But of enterprises that don't yet use MPLS VPNs but plan to, 67% plan to also deploy IDS. It might be concluded that network-based security services are an MPLS-usage driver for some organizations.

Figure 2. Network-Based Security Service Adoption

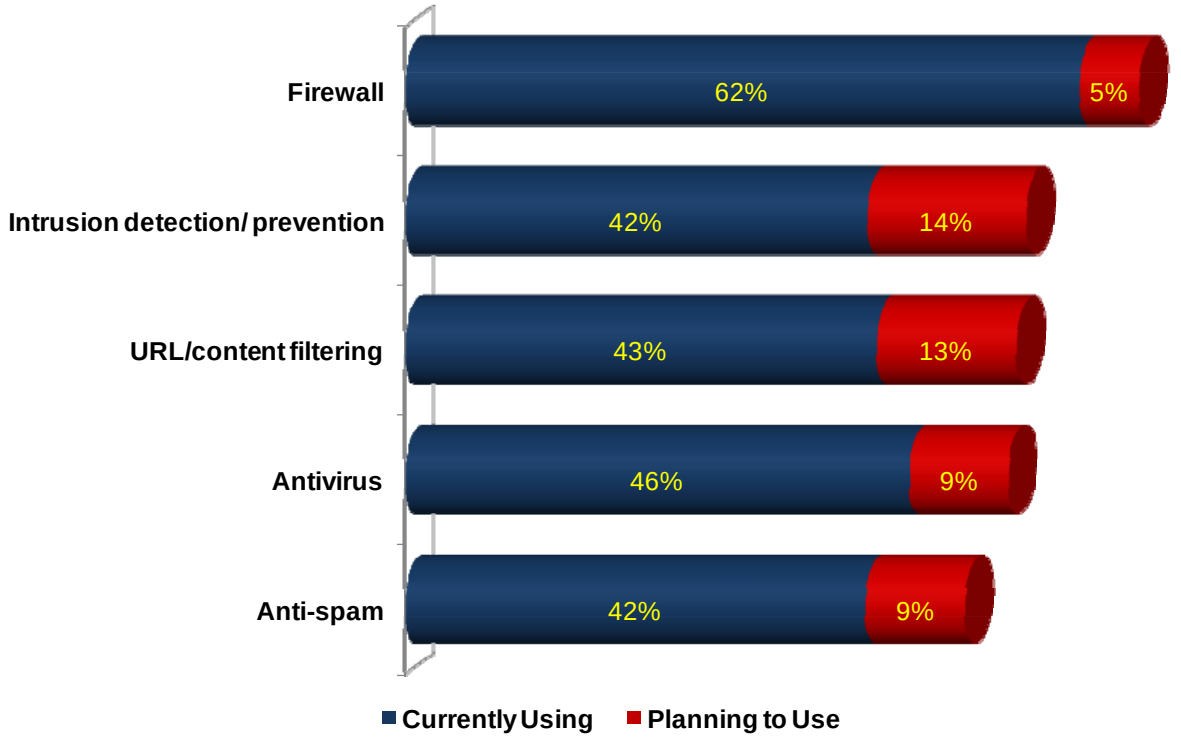

Current MPLS VPN customers' usage and planned usage for network-based security services.

Conclusion

Collaborative applications ranging from unified messaging to high-end telepresence are finding their way onto MPLS VPNs. Use of these applications is being partially fueled by the expectations set by mobile and social networks, which allow dynamic communications and file sharing and are spilling over from the consumer market to the business environment. In addition, enterprises have an economic need to find less expensive and more streamlined ways to build customer and business partner relationships and streamline productivity. This said, the broadband networks with the required bandwidth, QoS, performance management and privacy capabilities are available to support them.

MPLS VPNs, in particular, offer strong support for intranet collaboration because of their bandwidth and low-latency characteristics, including QoS packet prioritization and the ability to route directly from site to site. In addition, MPLS VPNs offer inherent privacy using VRF partitioning. A host of network-based security service add-ons are available for other layers of protection against malware, spyware and DoS attacks that might come from MPLS VPN gateway connections to the public Internet.

This report series was made possible, in part, by AT&T.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward-looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

The primary author of this Webtorials State-of-the-Market report is Joanie Wexler, an independent technology writer and analyst based in Silicon Valley.

ABOUT THIS WEBTORIALS REPORT SERIES

This report is Part 2 of a three-part series. The first report discusses the MPLS VPN survey's general key findings. In Part 3, a Webcast, lead analysts Joanie Wexler and Steven Taylor discuss where enterprises and industry perspectives matched up, where they didn't, and the possible reasons why.

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Professional Opinions Disclaimer

All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.