

Guaranteed performance

Could VPLS offer the answer to the seemingly inevitable future bandwidth crunch? Chris Werpy explores the options

With any communications network, the most common demand from multinational enterprises is for a reliable, secure and cost-effective communication channel between their globally dispersed offices, which requires guaranteed end-to-end bandwidth performance. However, bandwidth is under growing pressure from the increasing popularity of multimedia communications and converged voice, video and data applications, such as VoIP and video conferencing. As a result, the threat of traffic bottlenecks occurring between LANs is looming and corporations are looking for safe and guaranteed LAN to LAN connectivity that is scalable to meet whatever future bandwidth they may require.

Although carriers and service providers have been offering VPN services based on traditional TDM, Frame Relay, and ATM for some time now, the cost of operating separate networks to provide these services, coupled with the greater bandwidth consumption pressures, is forcing them to move to more cost-effective technologies: namely IP and MPLS.

Enter global virtual private LAN service (VPLS) into the networking spotlight. VPLS is a point-to-multipoint Ethernet-based transport service that allows businesses to securely extend their LAN throughout the entire WAN. VPLS benefits from the scalability and reliability of an MPLS core, with no legacy Frame Relay or ATM networks to integrate, and access to the existing network infrastructure and equipment. It scales well to national or international domains while preserving the quality of service (QoS) guarantees, with the added privacy and reliability of a Layer 2, carrier-class service.

The question is; are there any alternative solutions available that can compete with VPLS, such as Private IP, point-to-point solutions (such as Virtual Leased Line), Ethernet-in-Ethernet, L2TP and Border Gateway Protocol (BGP)/MPLS VPNs? The simple answer is "No". Simplicity and transparency is the name of the game for VPLS.



VPLS lets customers maintain control of their networks while allowing them to order on-the-fly bandwidth increments for multiple sites, instead of being constrained by the traditional legacy services. Configuration is also very straightforward — only the peer PE routers for a VPLS instance need to be specified. VPLS uses edge routers that can learn, bridge and replicate on a VPN basis. These routers can be connected by a full mesh of tunnels, enabling any-to-any connectivity. Customers can either use routers or switches with a VPLS solution, as opposed to Private IP. VPLS always offers Ethernet port handoff (customer demarcation) between the carrier and the customer router or a simple LAN switch allowing higher bandwidth service at a lower cost of deployment. Unlike IP VPN, where the customer hand-off can range from Ethernet, frame relay, or IP over TDM, with VPLS the customer hand-off to the WAN is always Ethernet. VPLS is also access technology-agnostic. The list of advantages is substantial.

BGP/MPLS VPNs (also known as 2547bis VPNs), on the other hand, require in-depth knowledge of routing protocols. As the number of instances increase, service provision-

VPLS lets customers maintain control of their networks while allowing them to order on-the-fly bandwidth increments for multiple sites

ing systems are often recommended in both cases to ease the burden on the administrator, particularly for Layer 3 VPNs. Layer 2 VPNs also enjoy a clear separation between the customer's network and the provider's network — a fact that has contributed heavily to its increasing popularity. Each customer is still free to run any routing protocol that the customer chooses, and that choice is transparent to the provider. Layer 3 VPNs are geared towards transport of IP traffic only. Although IP is nearly ubiquitous, there could be niche applications that require IPX, AppleTalk or other non-IP protocols. VPLS solutions support both IP and non-IP traffic. One significant security and performance advantage is that there is no IP interaction at the connection between the provider edge and the customers' devices.

Another differentiator is that VPLS offers greater flexibility and cost reductions, by putting the customer in control of the network and the negation of equipment upgrading requirements. End users have the flexibility to allocate different bandwidths at different sites, with the bandwidth varying from site to site by as much as 1 Mbps (for example at a low traffic-generating sales site) to Gig-E (which could be needed for the company's headquarters and/or data centre). Furthermore, as customers increase the bandwidth, there is no need to buy new cards for the existing CPE. I estimate that customers with a 50-site network can save up to 20 per cent in networking costs by moving over to VPLS.

VPLS solutions also score highly in the areas of compatibility and scalability. They are transparent to higher layer protocols, so that any type of traffic can be transported and tunneled seamlessly. VPLS auto-discovery and service provisioning simplifies the addition of new sites, without requiring reconfiguration at existing sites.

The most effective VPLS offerings are delivered using Ethernet connectivity, in the form of VLANs. These VLANs can be provisioned across TDM connections (E1, T1, E3, T3, etc) when native Ethernet is not available.

Dan O'Connell, research director for Gartner is a staunch supporter of VPLS, stating recently that "VPLS is a major new growth area for Ethernet. Customers are already very familiar with Ethernet in their local area networks. Extending to the wide area network is a natural progression, especially for those business and government customers seeking a clear IP migration path to enable

convergence of their multiple legacy networks."

So with VPLS displaying so many powerful capabilities, it would be surprising to imagine any circumstances where VPLS would be less than optimal. However, there is one such application and it's called multicast. Unlike Ethernet networks where there is native support for multicast traffic, VPLS requires the replication of such packets to each PE over each pseudo-wire, in order for multicast packets to reach all PE routers in that VPLS instance. The problem is further exacerbated in metro networks, where ring-based physical topologies are often deployed. Clearly, this replication is very expensive, causes wastage of bandwidth and is applicable at best when multicast traffic is expected to be a small proportion of overall traffic needs. Alternative solutions that the industry is researching include the establishment of shared trees within the VPLS domain, but this research has a long way to go.

Despite this caveat, VPLS is gaining momentum. Maria Zeppetella, Senior Analyst in Business Communications Services at Frost & Sullivan, agrees with this trend, but makes the point that most carriers are not planning to shut down legacy networks, as they still obtain a steady, albeit shrinking, revenue stream from them. Sprint being one exception, however, as it has announced it will shut down its legacy networks in 2009 and will have full migration to IP by then.

As for what the future holds for VPLS, I believe it will become the biggest network solution adoption of 2008 and 2009 for globally dispersed enterprises. Service providers will look to enhance their offerings for the early adopters and will introduce more customer network control applications and features. End users are always looking to simplify their network connections, while optimising transport effectiveness and keeping costs low. For businesses that are globally distributed and want to extend the benefits of the lower costs and simplicity of Ethernet technology throughout their entire network, global VPLS based on IP/MPLS is the solution of choice.

Chris Werpy is Director of Sales Engineering at Masergy, and can be contacted via tel: +1 (866) 588-5885; e-mail: chris.werpy@masergy.com www.masergy.com