VPLS vs. MPLS

You've probably heard of Multi-Protocol Label Switching (MPLS), but Virtual Private LAN Service (VPLS) may be new to you. It is now possible to consider either or both of these services for VoIP calls among sites. VPLS has much in common with a Layer 2 VPN in its implementation and configuration. So which is better? Find out if one of these technologies may be a better fit for your organization.

I always wanted to create a headline without words. This is my chance. You have heard of Multi-Protocol Label Switching (MPLS), but Virtual Private LAN Service (VPLS) may be new to you. MPLS has been available for a few years. VPLS was announced by Verizon in March 2007. AT&T and Qwest are also beginning to offer VPLS, which is a new service for both the service provider and enterprise, so the experience level is low. It is now possible to consider either or both of these services for VoIP calls among sites.

MPLS is a network function that is commonly a service, not a technology that an enterprise would install on their own routers. MPLS delivers a mechanism whereby IP networks can define virtual circuit services in a meshed connectivity that also improves security. The enterprise can assign traffic to multiple levels (4 or 5) of QoS. The enterprise can then request different performance guarantees regarding network latency (delay), jitter and packet loss. The enterprise connects to the MPLS service through local access lines (T1 or Frame Relay).

An MPLS-based private IP Virtual Private Network (VPN), also known as a network-based VPN, generally uses a service provider’s Ethernet as the backbone of the VPN. MPLS is a solution that allows the service provider’s core routers to transition from Layer 2-based connectivity to Layer 3. Additionally, the carrier’s core routers, not the end-user’s Customer Premises Equipment (CPE), create the VPN. For most service providers, there is no need to provision new circuits, if the enterprise is already their customer, and routers may not need to be replaced.

MPLS offers a fully-meshed architecture (meaning all sites can communicate directly with any other site without having to run through a hub/host location first), which has two key benefits for most enterprises. The meshed architecture improves site-to-site performance and imposes fewer burdens on remote locations. Network meshing and the addition of subsequent nodes are automatic functions of “connection-less” technology.

VPLS

VPLS is an Ethernet based service that looks like a Layer 2 VPN. VPLS supports geographically distributed Ethernet LANs. It uses MPLS as the transport/backbone network to carry the packets, but the interface at the enterprise faces what looks like a regional or national Ethernet. The enterprise’s locations appear to be on the same Ethernet LAN even though the packets traverse the service provider’s MPLS network.

VPLS has much in common with a Layer 2 VPN in its implementation and configuration. The customer CPE, a router or LAN switch, delivers the packet to the service provider’s provider edge (PE) device. The PE is a router facing the MPLS network. The MPLS network delivers the packet to the remote PE and then to the customer’s CPE. The difference is that the VPLS packet can be delivered in a point-to-point or multipoint fashion. The Layer 2 VPN is restricted to point-to-point connections. This means that VPLS can also broadcast over the VPLS network.

VPLS presents an Ethernet interface, simplifying the LAN/WAN boundary for service providers and customers. This enables fast and flexible service provisioning, because the service bandwidth is not tied to the physical interface. A 100-Mbps interface can support a Service Level Agreement (SLA) from 1 Mbps to 100 Mbps of customer traffic, commonly in increments of 1 Mbps.

Since VPLS uses the MAC address to locate the other endpoint, the IP address is not used. Therefore the enterprise should consider RTP header compression. This header compression can be implemented in the enterprise router facing VPLS. The advantage is a 20% to 60% reduction in voice bandwidth requirements.

So which is better for VoIP, MPLS or VPLS? As usual, it depends. VPLS is easy to implement with inexpensive switches. It has a wide range of bandwidth-on-demand options. VPLS is attractive for multiple locations in a metropolitan area. VPLS, however, is only good for tens of sites and hundreds of endpoints. If VPLS is used for a modest number of regional inter-site connections, then it appears to be good solution. The best approach is to interconnect routers at each enterprise site. Do not attach hundreds to thousands of IP phones directly over this service.

The initial startup uses a flooding transmission to all locations, just like Ethernet, to find the other devices. This works well for a few sites but does not scale well. The flooding consumes bandwidth on the network WHICH IS NOT MITIGATED by QoS. Buying faster bandwidth does not help.

If your remote offices use the Internet, not VPLS, to connect to the regional sites, it’s OK. If, however, you need to scale to hundreds of sites, this is not the technology. Select MPLS for a larger number of remote sites.

The VPLS SLA encompasses on- and off-net availability, MTTR (mean time to repair), round trip delay, packet delivery ratios and jitter limits. Ask the vendors under what conditions the broadcast flooding will occur, such as startup, recovery, server switching, power loss, etc. It may not be an issue, but the enterprise should know the influence flooding has under different conditions that will cause temporary congestion and performance degradation. Another question is, "Do the SLAs hold during flooding, or is flooding an exception?" Will the enterprise’s sites be backhauled to the VPLS service, or will they connect directly to VPLS? If backhauled, where does the SLA start?

Both VPLS and MPLS will support VoIP effectively. They each offer a different interface to the enterprise. VPLS is simpler but MPLS can grow bigger