

Virtual Backplanes and How They Affect Subnets

Edwin Hoffman CDO

OVERVIEW

The ability to create networks where a subnet can exist on both sides of a wide area network (WAN) connection has been a goal of many network engineers for quite some time. This has become more important in recent years because of latency-sensitive applications such as Voice, Video, and Storage over IP. Other applications now exist that are becoming latency-sensitive such as stock and foreign exchange trading applications.

LATENCY EFFECTS

Latency is caused in many ways. It never is just network transport that is the sole creator of network latency. For example, in VoIP transport almost 90% of the allowable latency is caused by the encoding and decoding of analogue voice sounds into digital data. Most often the compression used to reduce the size of a voice packet is the prime creator of latency. Latency also has a real cost in a network. If an existing network is to be upgraded to include VoIP, any high levels of latency need to be removed or bypassed, which increases the cost overall network, increases the cost of implementation, and finally increases the cost of maintenance. So although latency in a network is almost never the main cause of the overall latency for a stream of data, it is almost always the most variable and the most avoidable component.

As an example with voice, where we want less than 100–150 ms of latency in an end-to-end conversation (customers start noticing problems from 150 ms to 250 ms delay), the codec itself can create 35 ms at each end making 70 ms of delay (latency), jitter buffering (trying to even out the packet flow so it sounds like real voice) can add as much as 100 ms total delay giving us a 170 ms delay before we even start moving across a network.

Given we have around 50–80 ms available for network transport, end-to-end, latency becomes important. In a Layer 2 network, where overall latencies are counted in microseconds, latency is rarely a problem. Although distance over copper cable or fiber cable can also induce a problem (1 ms per 100 km of standard fiber). Typically, Layer 2 networks are ideal for VoIP networks.

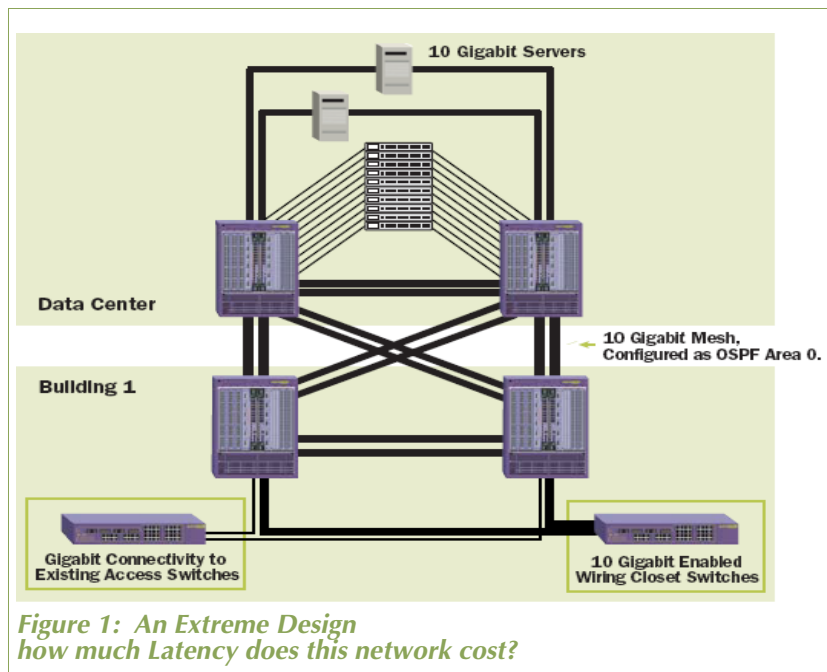


Figure 1: An Extreme Design
how much Latency does this network cost?

The problem arises when we have to traverse Layer 3 or routed networks, where latency through the router is often quite high. Routers were never designed to be low-latency devices, and improvements have not changed this fact in any real sense.

Layer 3 latency becomes a real problem when we design Layer 3 redundancy into networks, whether it is VRRP (standard), HSRP or EIGRP, ERSP (Extreme Networks), FRSP (Foundry Networks), or some other Layer 3 redundancy scheme. Latency is a major side effect of these redundancy choices. The moment we introduce these redundancy options, the cost rises for both initial implementation and for continued maintenance.

SUBNET PROBLEMS

Why would anyone want to create a subnet that existed on both sides of a WAN connection? There is not any routing between sides for that subnet, so costs are lower because this is a much simpler system. This would be the ideal low-latency design if it was possible.

How much latency and money does the network cost in Figure 2? In addition this double movement of data over an interconnect link adds both of the transit times (1ms/100 km fiber), and because it is increasing the bandwidth used, any delays caused by congestion, which can be high.

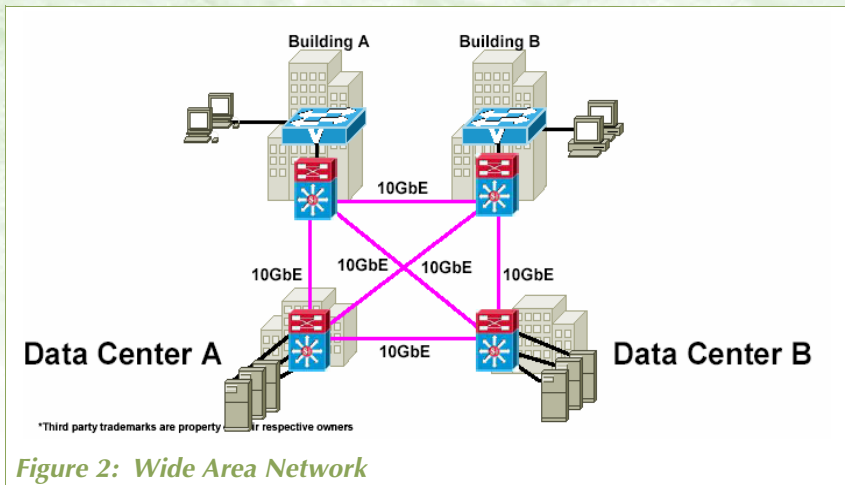


Figure 2: Wide Area Network

Reference the Tolly Group Test 201118 that shows how a switch can go from 15 microseconds to 1295 microseconds (1.3 seconds) of latency just because congestion was present.

Resilient Packet Ring (RPR) and Q-in-Q (VMAN) or VLAN stacking were introduced to try and create this “holy grail” and allow a VLAN to exist on multiple sites, which were connected using one of these technologies, but there is a problem. When any of these “attempts” to create a transportable VLAN are made, there are major issues that these technologies cannot address.

When a VLAN is created, the switch it was created on “owns” the VLAN. When we transport a VLAN over to another switch, we have to “trunk” the VLAN over using 802.1Q and 801.p tagging so that the remote switch can “join” the VLAN.

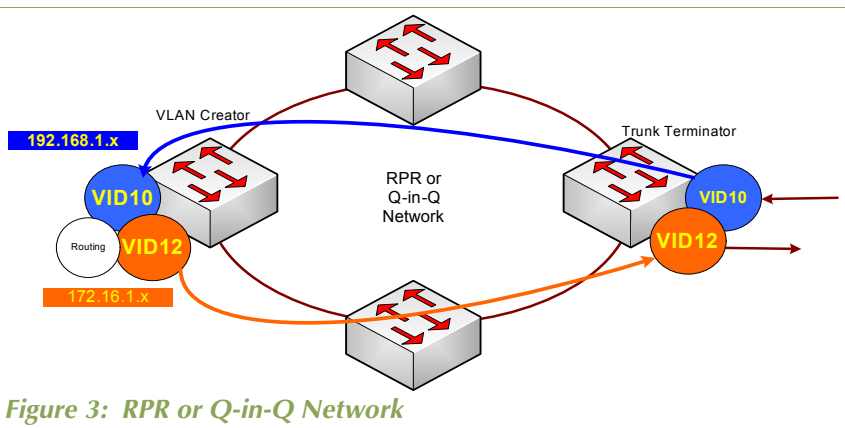


Figure 3: RPR or Q-in-Q Network

The problem with this structure is that when we decide to route between VLANs, we find that the only switch that can route the VLAN is the “owner” of the VLAN (the creator), so we have a situation that expensive bandwidth is wasted because we have to transport data “back” to the VLAN creator to be routed, and perhaps over the same link, to the egress point.

If we truly had VLANs that existed on both sides of the network, then the routing function would occur at both sides.

RAPTOR NETWORKS SOLUTION

With Raptor Adaptive Switch Technology (RAST™) from Raptor Networks, the VLAN does truly exist on both sides of the network. The subnet exists on both sides of the network, and routing requirements are carried out locally without wasted bandwidth. More importantly, it does not increase latency. Imagine what happens to latency when we have to transport data from side to side just to route it!

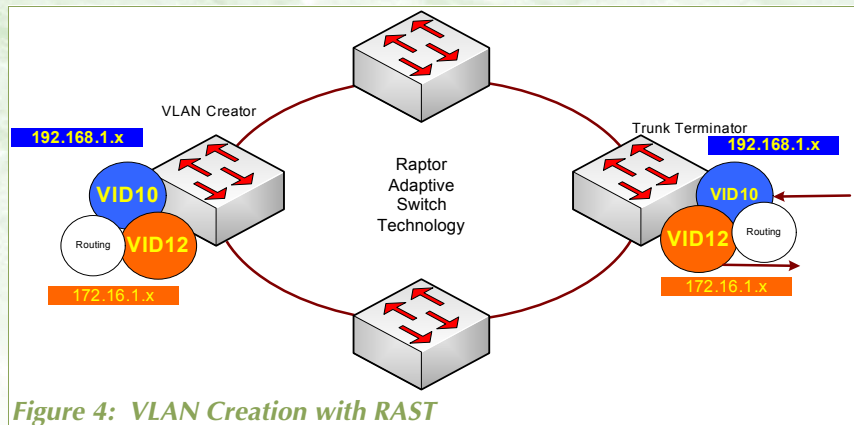


Figure 4: VLAN Creation with RAST

SUMMARY

Latency is a killer of certain types of application support. Often the fix is so expensive that it does not make financial sense, and the new application is shelved. When this shelving becomes excessive, companies cannot compete. As redundancy is introduced using Layer 3 industry standards, then latency is also greatly affected. All existing Layer 2 WAN methods or redundancy methods waste bandwidth because of the way VLANs operate.

Ether-Raptor with RAST allows the user to create VLANs and therefore IP subnets, on multiple sites concurrently, pushing back the point at which routing actually takes place and allowing multiple points at which routing decisions are taken on the same subnet. Now administrators can create campus designs, datacenter interconnects, or even just redundant backbones where the backbone is truly redundant and Layer-2-switched, with affordable, high-performance, 10 Gbps RAST links. This allows the heretofore impossible ability of creating subnets that exist on multiple sites.

Corporate Headquarters: 1241 E. Dyer Road, Suite 150 Santa Ana, CA 92705

Phone: 949-623-9300 / Fax: 949-623-9400 / Web: www.raptor-networks.com / E-mail: info@raptor-networks.com

Raptor Networks Technology, Inc. reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Raptor Networks Technology, Inc. is believed to be accurate and reliable. However, Raptor Networks Technology, Inc. does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Raptor Networks Technology, Inc. is a registered trademark and RAST is a trademark of Raptor Networks Technology, Inc. All other trademarks are the property of their respective owners.

CD-WP800 03/09/05