SearchEnterpriseWAN.com

Clarifying WAN virtualization, WAN aggregation and WAN optimization

Talari Networks, a pioneer in the WAN virtualization industry, is often lumped together with WAN optimization solutions. Does this mean the technologies are the same? And how does WAN virtualization compare with WAN aggregation?

To understand the nuanced differences among these breakthrough WAN technologies, SearchEnterpriseWAN.com spoke with Andrew Gottlieb, founder and board member of Talari Networks. In this Q&A, Gottlieb explains how WAN virtualization contrasts with WAN aggregation and WAN optimization.

What is WAN virtualization?

Andrew Gottlieb: WAN virtualization is a two-ended technology -- similar to the way [symmetrical] WAN optimization is deployed -- that enables network managers to use multiple WAN connections to augment or replace individual private WAN connections; [those connections can] be existing private WANs, such as MPLS, or Internet WAN links, like DSL, cable, fiber, Metro Ethernet, etc.

[WAN virtualization performs] bandwidth aggregation, allowing all links to be used almost all of the time. [It also] typically has loss-mitigation capabilities to minimize the application performance impact of packet loss, latency or jitter. A well-designed WAN virtualization solution does dynamic, real-time traffic engineering, reacting subsecond to not only link failures but also congestion-related network problems. This enables enterprises to build a WAN with much higher bandwidth for a much lower cost than a single-vendor WAN based entirely on MPLS.

How does WAN virtualization differ from WAN aggregation?

Gottlieb: There is confusion here both because the term WAN aggregation -- or its equivalent multilink WAN aggregation -- is used sometimes for single-ended technologies and sometimes for dual-ended technologies, and because even for dual-ended technologies WAN aggregation can be used to describe simpler, special cases of what WAN virtualization delivers.

WAN virtualization is unrelated to the single-ended definition of WAN aggregation, and is a superset of what's historically been delivered by products which do simpler WAN aggregation, most specifically in delivering excellent application performance and performance predictability in the face of network congestion problems that occur on shared networks like MPLS and especially the public Internet.

Note that all WAN aggregation and WAN virtualization technologies will correctly handle reliability when a network connection fails completely -- like when a WAN link dies, or a router fails. The more complex issue separating simpler WAN aggregation from WAN virtualization is handling connections with dissimilar latency, jitter and packet loss and, in particular, what happens in the face of network congestion-based changes in loss, latency and jitter.

What's an example of how WAN virtualization technology behaves in contrast to WAN aggregation?

Gottlieb: WAN routers supporting multiple WAN links, and using backup connections, are in some sense doing WAN aggregation, especially if you turn on balancing protocols like OSPF equal-cost routing, but no one would really describe this as WAN virtualization. Any given flow is steered to a given route using some predetermined method. Other than in the case of link failure, all flows of that type will stay on that WAN connection.

FatPipe was perhaps the first company to deliver WAN aggregation technology focused on supporting WAN links with dissimilar loss [or] latency characteristics, like multiple public Internet connections. While it handles link failures very well and, in fact, has had some success in the small- and medium-business market, because it always stripes packets from every flow across all available WAN connections, it is almost guaranteed to cause application performance predictability problems whenever there is congestion on any WAN connection. For this and other reasons [it] has never really caught on in the large enterprise market. That said, whether or not one describes what FatPipe does as WAN virtualization is perhaps a matter of degree.

WAN virtualization, from companies like Talari and Ipanema, goes behind traditional WAN aggregation in that it not only allows the use of multiple dissimilar WAN connections, but delivers excellent application performance and performance predictability even in the face of meaningful loss/latency/jitter events on the WAN. And [WAN virtualization] determines where to put traffic based on the current, real-time conditions of each network connection.

The best WAN virtualization implementations will allow you to use all available bandwidth even for just a single large flow -- like a file transfer or storage-backup application -- and can even move application traffic from one connection to another in real time in the face of network congestion, thus ensuring both maximum bandwidth utilization and predictable application performance.

It's the delivery of this predictable application performance in the face of an unpredictable network that enables enterprises to take advantage of inexpensive public Internet connections, including broadband connections which will see proportionally more episodes of network congestion than TDM links, and so makes WAN virtualization a much more comprehensive and valuable technology than simpler WAN aggregation technologies.

WAN virtualization encompasses WAN aggregation, but delivers more in terms of reliability and performance predictability, especially in the general case of combining multiple dissimilar WAN connections.

How does WAN virtualization compare with WAN optimization? Are the two technologies complementary?

Gottlieb: WAN optimization technology, which has been widely deployed for the benefits it brings to server centralization [or] data center consolidation projects, combines disk-based data de-duplication, application-specific help for Microsoft's CIFS protocol for file access, and some other TCP optimization techniques to save network bandwidth. [It] especially provide[s] application acceleration, which enables centralized server environments to deliver acceptable application performance.

WAN virtualization is highly complementary to WAN optimization. The majority of WAN virtualization deployments today add value on top of existing WAN optimization deployments. The bandwidth aggregation and loss-mitigation capabilities of WAN virtualization deliver further benefits to WAN optimization deployments, providing better end-user performance for first-time data transfers, and reliable voice and videoconferencing, as well as more predictable performance for all applications when any part of the private WAN becomes congested.

Where WAN optimization sees a reliable but expensive private WAN and seeks to squeeze every last bit of performance it can from the network, WAN virtualization sees lots of inexpensive, Internet-based bits out there, and seeks to make them reliable enough to be business quality. WAN optimization generally seeks to optimize each of many applications for the network, while WAN virtualization takes the approach of optimizing the network fabric for all applications.

For more information, listen to this podcast to understand WAN virtualization deployment and maintenance.

13 Apr 2012