

Architecting an SD-WAN with Quality of Service Over Broadband

Architecting an SD-WAN with Quality of Service Over Broadband

Introduction: Building a Better WAN	3
SD-WAN Building Blocks: Business Intent Overlays.....	3
The Elements of Tunnel Bonding	4
Link Bonding Policies	4
High Availability.....	5
High Quality.....	5
High Throughput	4
High Efficiency	5
Dynamic Path Control: Getting to Active/Active.....	5
Path Conditioning with Tunnel Bonding: Rock Solid Resilience	5
Traffic Shaping: A Two-Way Street	6
Putting It All Together	7

Introduction: Building a Better WAN


Internet connectivity is one of the cheapest and most widely available bandwidth options. However, when it comes to building a corporate wide area network (WAN), Internet connectivity is still not seen as a reliable medium for important business data. There are too many questions about whether broadband can support the quality and reliability enterprises have come to expect from their WAN. As a result of this skepticism regarding broadband quality, Internet links are often used as backup, where they largely sit idle.

What if the Internet could be made to be reliable? What if it could be secured and made to perform like a private line? What if technologies were available to enable active use of multiple sources of connectivity simultaneously AND deliver a level of quality of service (QoS) over broadband that increases reliability and security? Enterprises could rely on Internet bandwidth for essential applications. Instead of sitting idle, companies could utilize bandwidth they're already paying for. At the very least, they could augment their existing WAN and provide an easy way to scale into the future.

The concept of a software-defined WAN, or SD-WAN, has emerged as a new way to architect a WAN using multiple forms of connectivity. The challenge has been that while most SD-WAN providers address the ability to use multiple sources of connectivity, the need for performance improvements over broadband isn't addressed. Enter Silver Peak's Unity EdgeConnect platform. With EdgeConnect as the foundation for an SD-WAN, enterprises can improve the resiliency of their WAN, dynamically balance traffic across multiple paths and increase application performance.

SD-WAN Building Blocks: Business Intent Overlays

Network traffic traversing an EdgeConnect WAN deployment can be tuned for availability, quality, throughput or efficiency. This is accomplished on a per-application basis through the use of business intent overlays, also referred to as virtual network overlays. Multiple overlays can run on top of a singular physical infrastructure. Business intent overlays separate the functions of the network from the physical components of the network. Much like server virtualization converted physical servers to software instances, business intent overlays virtualize the network to increase scale, function and flexibility.




Business Intent Overlays create virtual networks and decouple the network function from the underlying network.

Business intent overlays make it possible to actively tune the network for a wide variety of needs and policies. The actual tuning of the network is made possible by EdgeConnect's ability to bond tunnels into a single, logical connection. Bonding links together, even tunnels, is not a new technology. However, when it is combined with other Silver Peak technologies like Dynamic Path Control (DPC), path conditioning, and dynamic rate control, true, end-to-end QoS can be realized across any type of bandwidth.

The Elements of Tunnel Bonding

In the following figure, there are three physical WAN links: one Internet link, an MPLS link, and a 4G/LTE link. (For clarity, we'll use "link" to refer to a physical circuits and the term "connection" to refer to a software-based tunnel.)



Bonded Tunnel Example

After building secure connections (AES encrypted IPSec) across each physical link, EdgeConnect can take two of these connections and bond them together into a single logical connection. In this example, the MPLS and the Internet links are bonded together. With link bonding, the resiliency of the connectivity is greatly enhanced. Should a single physical link fail, the logical bonded link remains up and packets are transmitted across the remaining active link(s) in the bond. In this example, the 4G/LTE link remains available outside the bonded condition to give additional capacity when needed.

Link bonding yields a highly resilient architecture that is simple to configure and support. Maintenance and changes can be executed on any single connection with no downtime, since the logical connection will still be up and operational.

Link Bonding Policies

Business intent overlays provide a means for tuning network parameters according to the application environment and desired network behavior. This feature of the overlay is configured by selecting one of four options, known as link bonding policies. The four link bonding policies are: High Availability, High Quality, High Throughput and High Efficiency. Link bonding policies combine the benefits of dynamic path control and path conditioning to simplify configuration and management of desired network behavior. Multiple business intent policies can be created, each with its own link bonding policy.

Link Bonding Policy	High Availability	High Quality	High Throughput	High Efficiency
	<ul style="list-style-type: none"> → Failover 0sec → Use Best Quality Path → Path Conditioning → BW Efficiency 50% 	<ul style="list-style-type: none"> → Failover <1sec → Use Best Quality Path → Path Conditioning → BW Efficiency >80% 	<ul style="list-style-type: none"> → Failover <1sec → Load Balance Paths → Path Conditioning → BW Efficiency >80% 	<ul style="list-style-type: none"> → Failover <1sec → Load Balance Paths → BW Efficiency 100%

Link Bonding Policy Options

High Availability

This policy is best for critical real-time services that cannot tolerate any interruption of service. Example applications would include call center VOIP, emergency services or critical VDI traffic. The goal of this policy is to maintain high quality throughput with high resiliency in the event of a brown-out or black-out. The High Availability setting also minimizes latency.

High Quality

This policy is best for typical real-time services, such as VoIP or video conferencing, online meetings, business-quality Skype, VDI traffic, etc. The goal of this option is to maintain high quality connectivity with sub-second failover.

High Throughput

This policy is best for any environment where maximum throughput is more important than quality. Examples of applications that could benefit from this option include data replication, backups, NFS, file transfers, etc. The goal of this option is to allow a few high-bandwidth applications to utilize the full-bonded bandwidth while still providing reasonable bandwidth efficiency

High Efficiency

This policy is best for network environments with many different kinds of applications. This is a good all-purpose setting for most TCP applications, with no FEC, no overhead, and just raw packets. The goal of this option is to allow a few high bandwidth applications to utilize the full bonded bandwidth balancing with 100% bandwidth efficiency


Dynamic Path Control: Getting to Active/Active

Dynamic Path Control (DPC) gives EdgeConnect the ability to intelligently route traffic over multiple available paths. Unlike traditional routing protocols, DPC monitors network conditions, in real-time, to make decisions about how to steer traffic across available paths. DPC will assess latency, packet loss, jitter and available throughput to choose the most appropriate WAN path for any application. Dynamic Path Control provides a viable means for implementing an active/active multiple link configuration. Once business intent policies are set for each overlay, DPC then dynamically directs traffic based on the defined requirements. In the event a link goes down, Silver Peak's DPC automatically fails over to the other active link in less than a second.

Path Conditioning with Tunnel Bonding: Rock Solid Resilience

A key technology that can increase path reliability is **path conditioning**. Path conditioning refers to two techniques - forward error correction and packet order correction – that negate the impact of packet loss and out of order packets across a WAN. Forward error correction (FEC) can reconstruct lost packets, thereby avoiding retransmissions and throughput degradation that plague IP protocols in the presence of loss. Packet Order Correction (POC) automatically detects and fixes packets that arrive at a destination in the wrong order. Again, this avoids costly retransmissions and throughput degradation. Overall, path conditioning increases the efficiency of lower quality links. When path conditioning and tunnel bonding are combined, network resiliency is greatly increased.

Let's look in more detail at how its underlying technologies--forward error correction and packet order correction--work in the context of bonded tunnels:



With path conditioning on EdgeConnect, lost packets are rebuilt using FEC. More precisely, EdgeConnect's forward error correction employs a technique called erasure coding to correct for lost data. Erasure coding is a time-tested paradigm that provides a sophisticated and efficient method for recovering lost data across the WAN. When used in conjunction with tunnel bonding, the error correction data can be sent down one connection (in this example, MPLS) and the user packets down another (Internet in this example).


With EdgeConnect, the level of protection is configurable. Each error correction packet contains data about all four data packets, but resiliency can be improved by configuring FEC for as much as a 1:1 ratio. A FEC ratio of 1:1 means one FEC packet is sent for every data packet. This is the most aggressive setting and is typically used in environments with significant packet loss, e.g. satellite links, military deployments, and long-distance, low-quality connections. With the right FEC configuration, EdgeConnect can recover from the loss of multiple packets and multiple FEC packets. This degree of resiliency across any WAN link is unparalleled.

Traffic Shaping: A Two-Way Street

Another important element of EdgeConnect is the traffic shaper engine. The shaper will shape traffic by allocating bandwidth as a percentage of the configured system bandwidth for each appliance. The shaper supports up to 10 traffic classes, with four (4) pre-defined classes for real-time, interactive, default and best effort.

EdgeConnect's shaper supports shaping of both inbound and outbound traffic. In the outbound direction, the shaper is can guarantee minimum and maximum bandwidth levels to applications, VLANs, subnets and overlays. Inbound shaping can be effective in all network environments, but it is mandatory across Internet links to prevent low priority traffic (e.g. streaming cat videos on YouTube) from overwhelming critical business traffic.

In conjunction with bi-directional shaping, a unique feature called Dynamic Rate Control (DRC) allows the EdgeConnect to automatically adjust the maximum bandwidth allocated to each tunnel (Tunnel Max Bandwidth) in order to avoid congestion across the WAN. In the example below, tunnels from sites A and B are auto-negotiated to a bandwidth of 155Mbps. However, since the Hub is connected via a 100Mbps link, the tunnel max bandwidth is auto-negotiated to 100Mbps outbound from the hub, as well as from site A and site B. Additionally, the inbound shaper is limits incoming tunnel throughput to 100Mbps, thus avoiding congestion in either direction.



EdgeConnect uses dynamic rate control and inbound shaping to avoid congestion.

Putting It All Together

Silver Peak's Unity EdgeConnect SD-WAN solution serves to securely connect users to their applications, traversing the WAN using the highest performing and most cost-effective means possible. Implementing business intent overlays that incorporate link bonding, path conditioning, dynamic path control, bi-directional shaping and dynamic rate control would be complex if these features were left to be set up independently. However, Silver Peak's SD-WAN platform makes this manageable. From a single central console, called Unity Orchestrator, network administrators don't have to spend hours tweaking and tuning, they simply use configuration wizards to give Orchestrator the basics (connections, bandwidth, topology, etc.). From there, the EdgeConnect goes to work to automatically configure secure network connections, appropriately according to each site's topology. From there, the overlays are built and configured with the appropriate policies.

In the end, making the Internet reliable for an enterprise WAN requires the seamless integration of many different technologies. From packet-based load balancing, to reliable bi-directional shaping, to path conditioning, to business intent overlays that virtualize network functions, EdgeConnect gives enterprises the unparalleled control over their WAN. Layer on top of all this, the option to add modular WAN optimization when and where it's needed (another unique feature of Silver Peak's SD-WAN solution called Unity Boost), and Silver Peak's EdgeConnect platform separates itself from the competition. Only Silver Peak delivers a complete SD-WAN solution that scales to meet the connectivity and performance needs of any enterprise.