SD-WAN SOLUTION GUIDE

bigleaf networks

velocloud

PRESENTING THE STRATEGIC AND ARCHITECTURAL FOCUS OF 2 INDUSTRY LEADERS IN SD-WAN TECHNOLOGY

mtrx.com

Contents

Introduction from the Author 1

Bigleaf Networks' Unique Architecture 2

VeloCloud's Networks Unique Architecture 5

Matrix Networks' Take 9

Introduction from the Author

A Changing World of Connectivity

Software Defined WAN or SD-WAN is an emerging networking technology that has evolved from Software Defined Networking (SDN). SD-WAN is designed to gather intelligence about your circuits, whether public, private, or a combination, and then make intelligent routing decisions based on the performance of each connection.

This is in contrast to traditional routing, which only gets its next hop information at layer 3 which is totally ignorant to the health of any given path or circuit. Traditional networks also often only have a single path. SD-WAN solutions try to always leverage multiple paths and then use the gathered intelligence to decide which path to use. The result is a brilliant new networking solution that improves the performance of Cloud applications like voice and video, while allowing for the replacement of expensive, slow legacy MPLS type private connectivity. Users of SD-WAN gain consistently stable, high performance access to critical applications, whether located on a server in a headquarters location or in the Cloud.

SD-WAN could still be considered an emerging technology, so the way in which each vendor approaches design is quite different. Some offerings are laser focused on displacement of the legacy MPLS networks; others are zeroed in on making connectivity to SaaS applications very reliable; still others offer some combination of these two principles in a hybrid solution. Understanding the way an SD-WAN solution operates is critical to identifying the topology that will fit into diverse business strategies. In this whitepaper we are focused on uncovering the unique architecture of Bigleaf and VeloCloud, industry leaders in this new space. Although both solutions are “SD-WAN,” their strategy for providing services is very different.

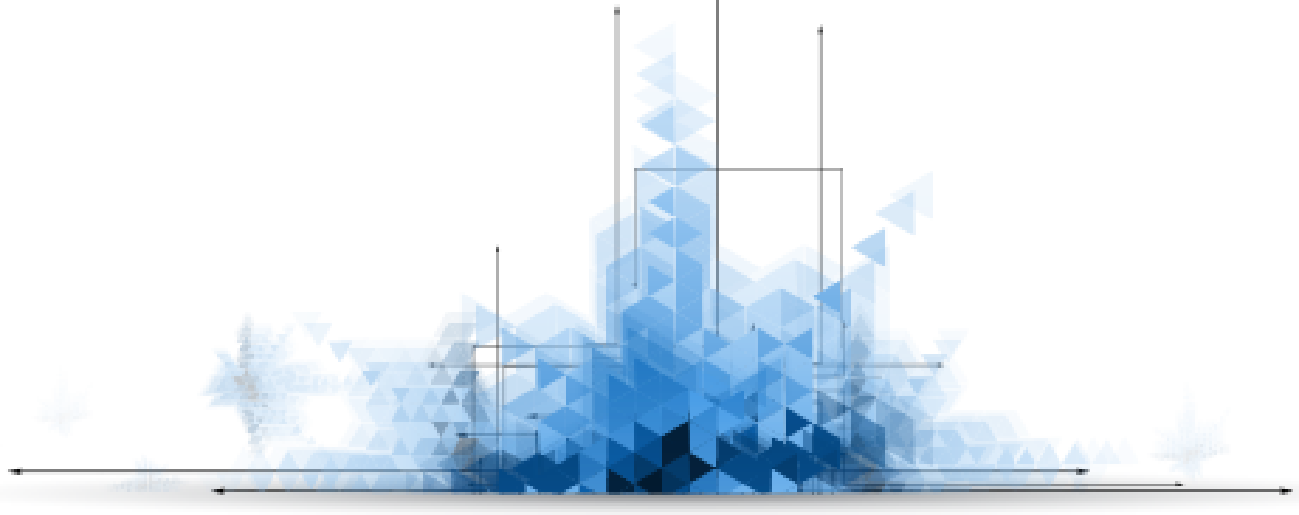

Introduction from Bigleaf

Bigleaf Cloud-first SD-WAN is the next generation of Internet optimization – based on the natural redundancy found in leaf veins. Bigleaf connects you through their plug-and-play router in your office, to Bigleaf's gateway clusters in the Cloud, providing end-to-end visibility and control. Bigleaf's Software Defined Wide Area Network (SD-WAN) technology ensures that your business-critical applications are prioritized across your Internet connections, and that all traffic flows are routed over the best connection in real-time.

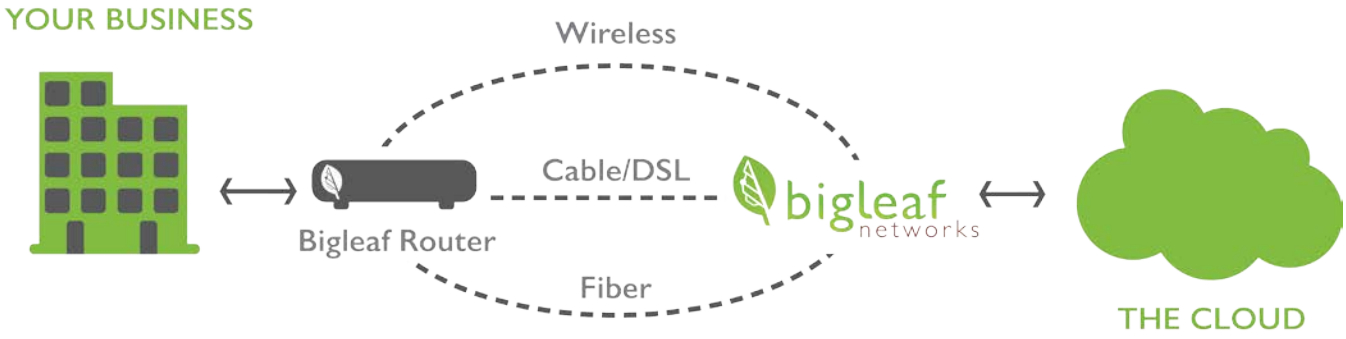

Bigleaf Networks' Unique Architecture

Bigleaf Networks has a unique approach to SD-WAN, self-described as "Cloud-first." Bigleaf's nationwide network includes a multi-carrier, fully-redundant backbone and 5 data centers around the US: Seattle, Los Angeles, Dallas, Chicago, and New York. Each of these data centers, where Bigleaf collocates its own dedicated infrastructure, has strategic peering connections with major Tier 1 Internet and SaaS providers, with the intention of getting as close as they can to the leading Cloud providers. This is the "Cloud-first" approach. With these peering relationships in place, they use SD-WAN technology to get your traffic to these data centers, and in turn, to SaaS providers, with far more intelligence than using a traditional Firewall and egressed traffic onto the public Internet.

The edge of the Bigleaf solution is a local appliance at each site used to terminate Internet circuits. The appliance then absorbs the addressing associated with these connections, normally assigned either via DHCP or static IP(s). A unique strategy of the solution is that this appliance is "in front" of the existing router/firewall, running in tandem.

Compare this to nearly every other SD-WAN solution, which wants to be the router/firewall. In the Bigleaf solution, you have the freedom to easily maintain any existing firewall. As the Bigleaf appliance absorbs your existing WAN IPs, they assign you an IP(s) out of their data centers. Their appliance creates transparent tunnels to one of their data centers over the Internet circuits, at which point all traffic egressing from your firewall emerges onto the Internet at the Bigleaf data centers.

The current model of the Bigleaf appliance can terminate up to 4 circuits and is best purposed with at least two diverse circuits. These circuits are simply treated as "paths" to their data centers, and any additional circuits terminated into the appliance gives sites more paths to the data centers. The appliance then sends packets 10 times per second over each path to the Bigleaf data centers. Using these packets, the appliance can determine the latency, packet loss, and jitter of each path. Using these metrics the solution intelligently routes traffic across the appropriate path, depending on the type of traffic.

When first signing up for the service, Bigleaf does need some basic information about the deployment location and the circuits it will terminate. The appliance is essentially “Zero Touch” in that by giving information up-front to Bigleaf, they pre-program the appliance, making it plug and play. An important part of provisioning is determining how many IP addresses are needed. The service always comes with at least 1 static IP. Additional IPs can be requested by filling out an IP justification form. If you have a BGP AS number, Bigleaf can announce it on your behalf and associate any IP blocks with your service. Using the physical location of where the appliance is going to be, the appliance is preprogrammed with the Bigleaf datacenter it will call home. This is determined based on geographic distance from the location to the closest data center, in order to achieve the lowest latency.

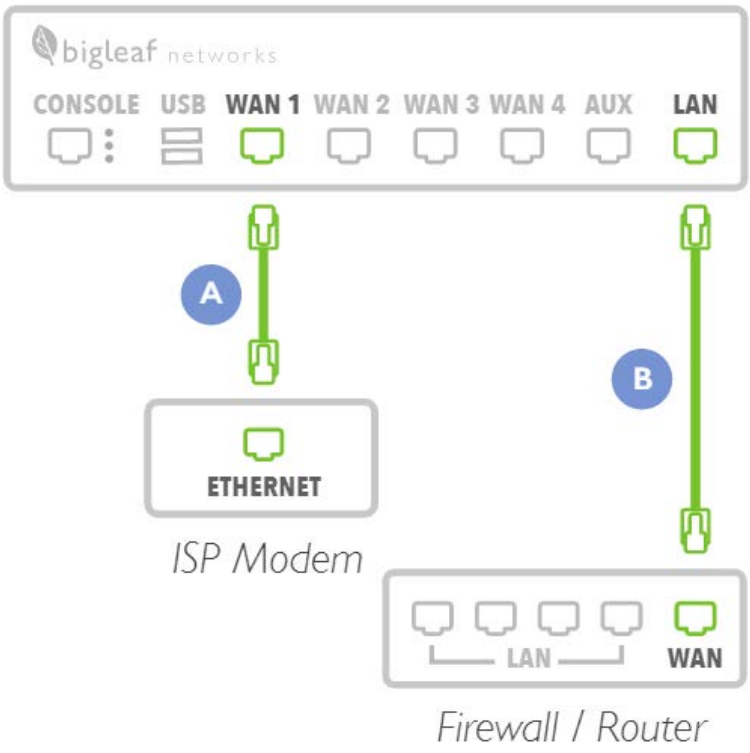

As mentioned, installing the Bigleaf involves putting it between the Internet connections and any existing firewall or router. From the firewall or router's point of view, you essentially go through the same process as if you were doing an ISP change. As part of this process you re-address the WAN side, and modify any existing NAT port-forwards that were in place. Forward and reverse DNS records also need to be considered if you are hosting any services on-prem. Another unique part of the solution is when you do host things on-prem, you still get the QoS overlay functionality for these services. This allows traffic in both directions, ingress and egress, to use the best path for the type of traffic being sent. If you do have on-prem VPN or RDP services that need to be always available, this is a major strength of the Bigleaf solution.

If there is an issue at one of the Bigleaf data centers, and their server-side infrastructure is unavailable, the IP block you are using is moved to another data center via BGP. The IP block is moved between data centers in under 5 seconds, however it takes about 30 seconds for that change to be advertised to the peering providers. The on-prem appliance is already preprogrammed with which data center the block is going to be moved to. Once the appliance is no longer getting those heart beats 10 times per second back from the currently connected data center, it then establishes a tunnel to the “backup” data center.

For even more redundancy, the Bigleaf appliances can be installed in a Highly Available (HA) configuration. This option is beneficial if HA firewalls are already in place, as it can give you multiple paths to each Bigleaf appliance. The HA solution involves 2 network switches, and 2 Bigleaf appliances with each Bigleaf plugged into both switches. ISP connection 1 is plugged into Network Switch 1, ISP connection 2 is plugged into Network Switch 2. The cleanest way to connect the firewalls into the mix is using LACP (link aggregation control protocol). With LACP you can plug both firewalls into both switches, which gives you maximum redundancy. If the firewalls don't support WAN LACP, you plug firewall 1 into Switch 1, and firewall 2 into Switch 2. Using a single firewall with LACP can be used and gives you protection if one of the network switches or one of the actual SD-WAN appliance fails. Using just a single firewall only plugged into Switch 1 doesn't provide much High Availability.

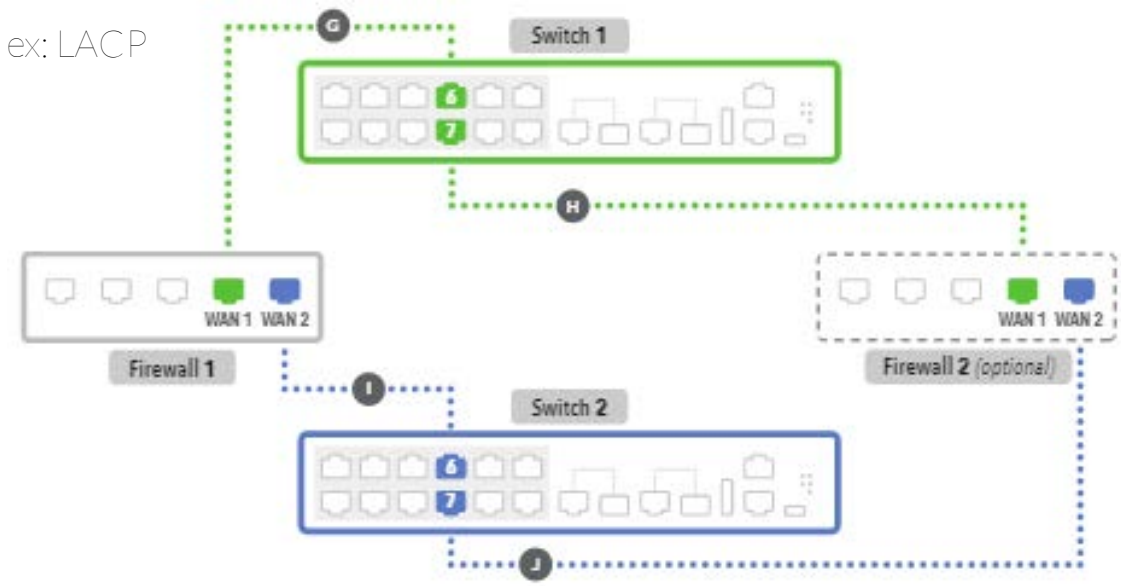

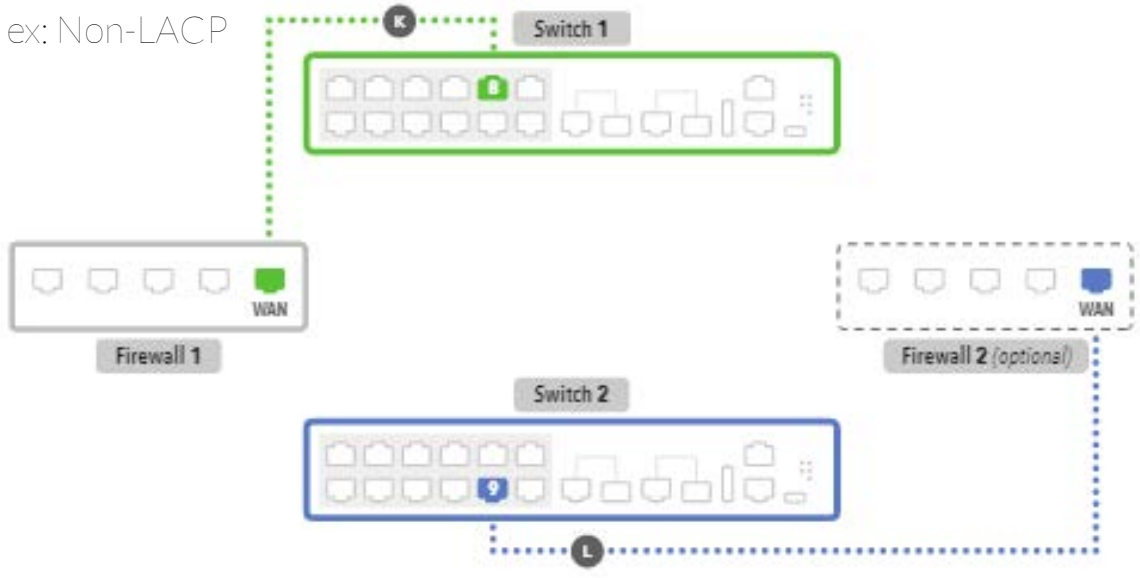

Introduction from VeloCloud

VeloCloud is the first to provide all three elements needed to achieve a Cloud-Delivered SD-WAN: a cloud network for enterprise-grade connection to cloud and enterprise applications, software-defined control and automation, and virtual services delivery.

VeloCloud is the only SD-WAN solution supporting data plane services in the cloud, in addition to on-premise deployments. This enables policy-based access to the cloud and data center applications. VeloCloud's SD-WAN leverages the economics of the cloud to offer a SaaS-like subscription price model to ease adoption and pay as you grow.

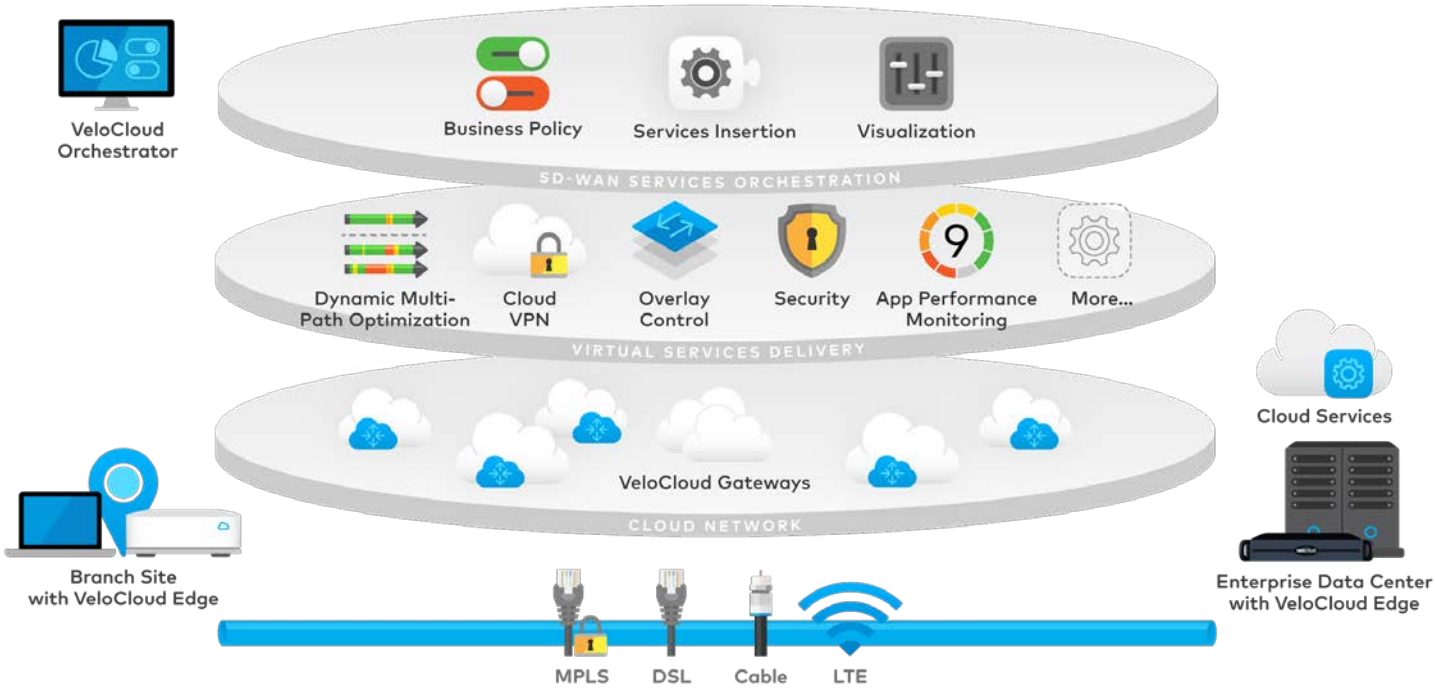

VeloCloud's Unique Architecture

VeloCloud has built an SD-WAN solution incorporating site-to-site and site-to-Cloud solutions. Their solution involves an on-prem appliance called a VeloCloud Edge, NFV software called a VeloCloud Gateway, and the VeloCloud Orchestrator. The VeloCloud Gateways exist in data centers and on Cloud provider infrastructure. They are strategically placed with SaaS providers in mind, with their promise to have the VeloCloud Gateways within 5ms of the major SaaS providers. The VeloCloud Edge then uses SD-WAN technology to connect to the Gateways, intelligently getting your traffic to the SaaS providers, and sites on the Internet at large. VeloCloud also has an interesting play in a multi-location WAN environment, as they can create tunnels between branch locations and a "hub" using multiple Internet circuits, or even a hybrid WAN using a private circuit and Internet circuit(s).

The VeloCloud Edge is the heart of the on-prem piece of the solution. There are currently 3 models of the appliance, with differing amounts of throughput and WAN/LAN ports. The smallest appliance supports 100mb of throughput with 3 WAN interfaces. The largest can handle 1 GB of throughput with 6 WAN interfaces. The WAN interfaces on all models are a mix of copper and fiber. The Edge is designed to have multiple paths to allow it to choose the best one for different types of traffic. These paths can be Internet circuits, private circuits, and even USB cellular modems, allowing for a combination that makes sense for each location. With this versatility, there are many ways to deploy an Edge into an existing environment. The Edge can be configured to be the NATing router for a site and just terminate Internet connections into it. When deploying into a private WAN environment (MPLS/ VPLS /EVPL), you simply set the next hop of the private WAN interface to the private WAN router and use a single Internet connection as the second path. If there is an existing security appliance at a location, there are a few options if it needs to stay in production. If the desire is to maintain the firewall as the security appliance, you can put it in routing mode and place it behind the Edge.

You can then use static routes, or a dynamic routing protocol like BGP or OSPF, to set up routes between the security appliance and the Edge. If the NATing router needs to maintain the private or Internet connection, you can 1:1 NAT the Edge (so it still has two paths) and utilize routing to get LAN traffic to pass through the Edge. In this configuration, you would want the Edge in routing mode, so the security appliance can still see the source addresses, allowing different polices to be applied. These are just a few ways an Edge can be inserted into an existing network.

For multi-location deployments, an Edge can create tunnels to other Edges for sharing on-prem resources, like VoIP or video calls between branch locations. To facilitate these tunnels between the Edges, the solution has built-in mesh VPN functionality. Edge appliances also have Wi-Fi functionality built in. While not as robust as a traditional AP, it does support 802.1x (PEAP). This can be ideal for a smaller branch location where there isn't a need for a full-blown Enterprise-grade access point.

Characteristic	Edge 500	Edge 1000
WAN Throughput	100Mbps with services enabled	1Gbps with services enabled
EXTERNAL INTERFACES		
WAN ports	2 x GbE, RJ-45, 1 x SFP	4 x GbE, RJ-45, 2 x SFP+
LAN ports	4 x GbE, RJ-45	4 x GbE, RJ-45
USB ports	4 x USB 2.0	2 x USB 2.0
Wireless Wi-Fi	802.11a/b/g/n Dual band (2.4/5GHz)	N/A

There is a simple local interface on the Edge router which can be accessed via HTTPS. This is only used for the initial setup, and local troubleshooting. Actual configuration is done via the second major component of the solution: the VeloCloud Orchestrator. The Orchestrator is hosted by VeloCloud and is also accessed via HTTPS. All associated Edges show up in the Orchestrator and can be managed from a single interface. Any changes made are then pushed out to the affected Edges. If the Edge and Orchestrator cannot communicate for any reason, the configuration of the Edge can no longer be changed, but it does continue to operate with the running configuration. In the Orchestrator there is a notion of Edge profiles; the profiles are then applied to Edges. This allows for mass updating of any Edges assigned to the profile, while still allowing for control of a single Edge by specifically editing the individual Edge. In addition to the "configuration" profiles, there are also "networks."

"Networks" allows you to set up the subnet ranges needed. When a new Edge is deployed, it chooses a subnet out of these ranges. In production you would generally overwrite this subnet at the individual Edge configuration, as there is a good chance the Edge is being deployed into an existing subnet. The Orchestrator includes incredibly detailed monitoring information. Health of the different paths is displayed as QoE (Quality of Experience) charts. Any latency, jitter, and packet loss are displayed for each path, then an overall chart shows what the experience was with SD-WAN multi-pathing. In addition to this chart, VeloCloud associates a "Quality Score" to each path, and then a total score of all the paths combined using the multi-pathing. The charts are very useful to see if there are any long-term issues on the circuits. There are also netflow-style source and destination reports for all traffic that was routed through the Edge. This information can be very useful for finding bandwidth hogs and for generally knowing which clients are accessing what.

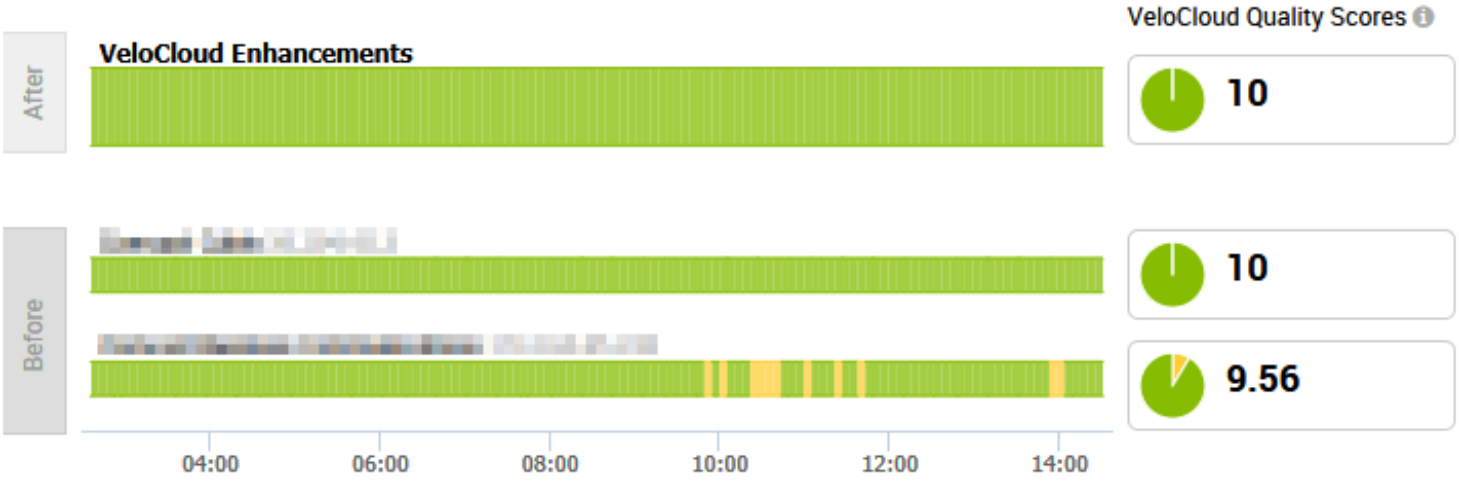

The third major component of the solution is the VeloCloud Gateways. When accessing the Internet through the Edge, by default it creates tunnels over any available path to the Gateways, and Internet traffic is egressed from the Gateways onto the public Internet. With every packet sent between the Edge and Gateway, the latency, jitter, packet loss, and bandwidth are calculated for each path. This information allows for the intelligent routing of traffic based on each path, creating a type of QoS overlay. If the Edge sees a real-time UDP application, it will send it over the path with the least latency and jitter. If there is a large TCP-based download, it might favor the path with the most bandwidth.

As the Edge identifies SaaS applications, it will try to create tunnels to the Gateway closest to those applications and send the SaaS traffic to that Gateway. The Edges can currently identify over 2,500 applications and websites. "Business Policy Rules" can be created to prioritize and route specific traffic over different paths. Rules can be matched on source/destination/application. Actions can include egressing traffic straight out one of the paths, tunneling that traffic to a Cloud Gateway, or even just backhauling the traffic to a hub. A hub is going to typically be a data center or co-lo. The VeloCloud solution is priced on ingress/egress to the Cloud Gateways, therefore organizations might not want to use that Gateway ingress/egress capacity for all traffic. So for an application like Netflix, you could just egress it directly out one of the Internet connections.

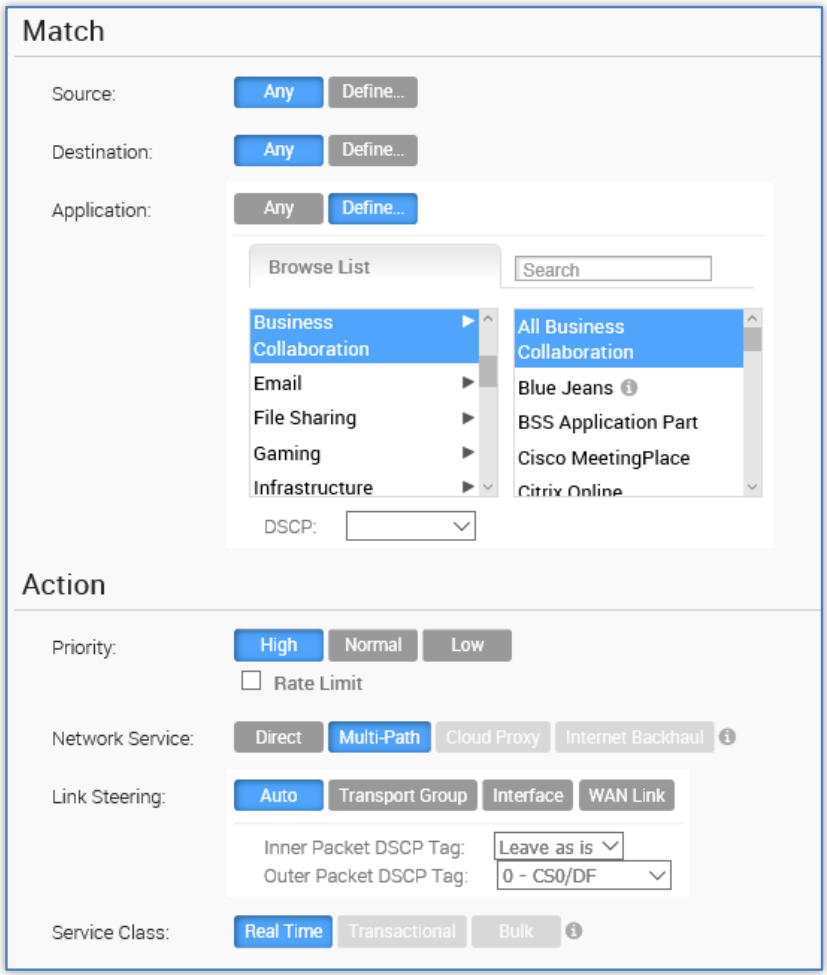

The Edge does have a basic firewall, with the advantage of also being able to match traffic based on application type (layer 7). The firewall administration interface is very similar to the Business Policy Rules for firewall matching, however the action will either deny or allow. For organizations used to a full-featured security appliance, the Edge doesn't have the more advanced features such as IPS/IDS, SSL Inspection, Anti-Virus, or Botnet detection. As mentioned, there are ways to use a security appliance in conjunction with an Edge. If you are going to use an Edge without a local firewall and need more advanced security services, VeloCloud does partner with Zscaler and Websense. Zscaler and Websense are highly rated Cloud-based security offerings. Access to these services is facilitated by VeloCloud, establishing an ISPEC tunnel to either of these providers from one of its Cloud Gateways. The local Edge then uses multi-pathing to send traffic to that Cloud Gateway, which is then routed over the IPSEC tunnel to a provider. At that point you can route traffic to either provider and take advantage of their more robust security services. This same technique of creating an IPSEC from one of the Cloud Gateways is also used if you have a site without an Edge. You might need this functionality to connect to some other 3rd party or if you have a co-lo/private cloud that doesn't have an Edge deployed. An Edge does have many of the expected NATing router features like port forwarding, 1:1 NAT, firewalling traffic routed between Interfaces/VLANs/subnets, and SNMP.

The Challenge of Choice…

Matrix Networks' Take

It is easy to see that SD-WAN brings in a whole new era of intelligent routing of traffic across any given path. This is really huge and can help companies of all sizes, even if they are going to continue to keep applications on-prem or are migrating to the Cloud. Choosing the best SD-WAN product for an environment, and deciding the best way to implement it, can be a challenge. Just because a solution is labeled SD-WAN does not necessarily mean it offers the critical features your business needs. **The term SD-WAN is just as ambiguous as the term "Cloud." What is the Cloud anyway?** There are many differences across all of the "Cloud" providers, and then there are even more ambiguous terms like private vs public when it comes to the "Cloud." SD-WAN is pretty much the same "Wild Wild West" environment in this regard. Keep your business healthy and do your homework, then work with trusted advisors (like Matrix Networks) to identify solutions that make the most sense for your specific network.

Want More Awesome Content?

Visit www.mtrx.com/whitepapers

This Solution Guide was written, edited, and published by Matrix Networks. The content of each section is intended to educate based on the experience and knowledge gained by Matrix Networks as a whole. The content is not intended to sway potential consideration for one vendor over another but rather to inform our audience from a third-party perspective in an otherwise partisan environment. For more educational articles and opinion pieces please visit www.mtrx.com.

About Matrix Networks

Matrix Networks provides IP Communications for the Growing Business – we work with organizations of all sizes to leverage the power of information technology to help organizations of all sizes achieve their IT and communication goals. Our solutions include Voice (Phone Systems), Wireless Internet (Wi-Fi), Video Collaboration, and an entire suite of Cloud Solutions ~ supported by MCaaS.

Matrix was founded in 1984 and as the communication industry has evolved, so have we; growing from a one man organization in N. Portland, Oregon to become one of the Northwest's longest lasting communication providers. At Matrix Networks, we are dedicated to providing you with a client experience that exceeds your expectations and delivers the industry's best solutions for Voice, Video, Wi-Fi, and Cloud – with Matrix Networks you'll have a true communication partnership with someone who understands your organization, cares about the direction and growth of your business, and will be there for you for years to come.

matrix™ NETWORKS