

4 PRACTICAL CONSIDERATIONS FOR SD-WAN DESIGN

AUTHORED BY

Steve Womer

Director, Solutions Engineering



INTRODUCTION

WHAT IS SD-WAN?



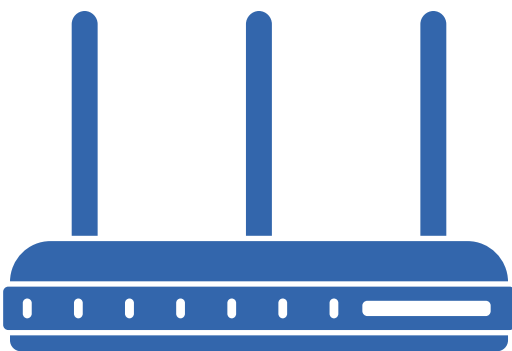
SD-WAN solutions provide a replacement for traditional WAN routers and are agnostic to WAN transport technologies. SD-WAN provides dynamic, policy-based, application path selection across multiple WAN connections and supports service chaining for additional services such as WAN optimization and firewalls.

- Gartner



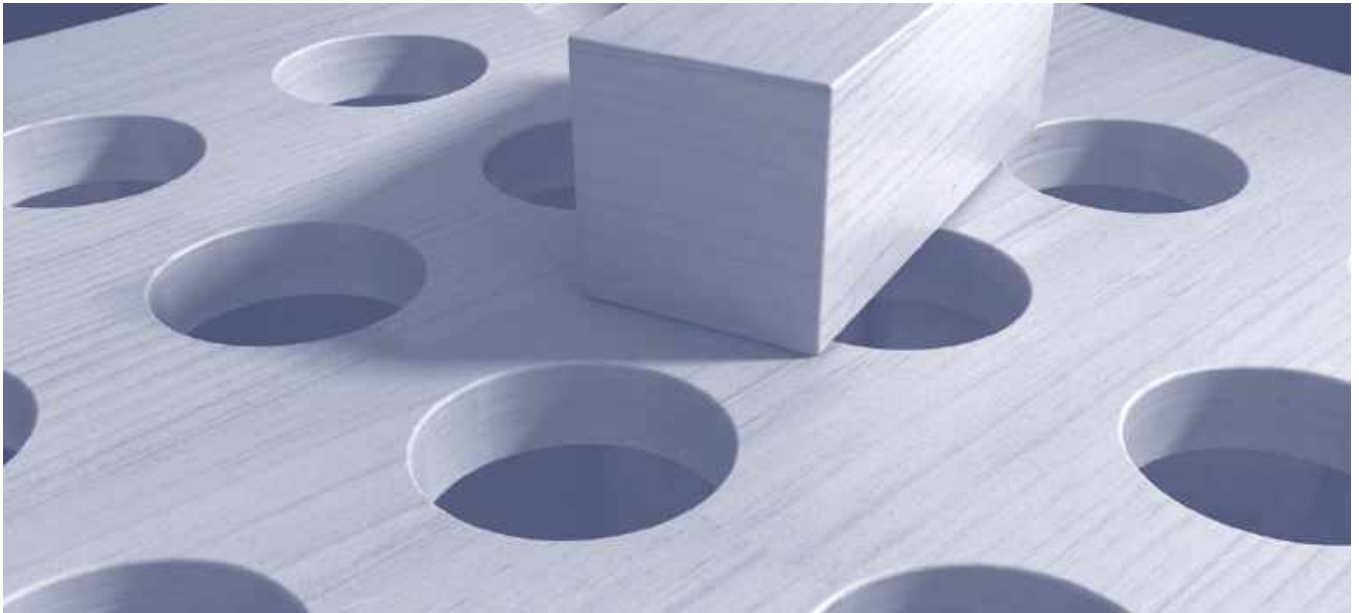
According to IDC, the Software Defined Wide Area Network (SD-WAN) market is expected to grow at 30.8% compound annual growth rate (CAGR) from 2018 to 2023 to reach \$5.25 billion in revenues. As enterprises seek to accelerate the adoption of cloud apps and drive greater efficiency and reliability from their network infrastructure, SD-WAN vendors are riding this wave of growth.

Using the definition by Gartner as a starting point, SD-WAN addresses **three fundamental considerations** in WAN networking:

**1****Initial Configuration and Deployment****2****Speed and Reliability****3****Network Visibility and Control**

The promise of SD-WAN is to enhance WAN performance and simplify WAN administration through software and automation. SD-WAN is an evolution of technologies that pre-existed the term "SD-WAN" but work extremely well as a unified architecture. Realizing the promised transformational benefits of SD-WAN is not a walk in the park. There are three fundamental challenges that enterprises need to watch for.

CHALLENGES WITH SD-WAN IMPLEMENTATION



Complexity and change management is often underestimated

SD-WAN solutions are not self learning (yet) and require a user to define what traffic is allowed on the network and how important that traffic is. One option for configuration is to set all traffic at equal priority and let the SD-WAN software do its best to distribute the load across the available bandwidth. That may be fine if you only have one or two key applications in your network, but for the majority of businesses that have a variety of applications that need to perform flawlessly, that strategy will fall short.

For an optimal end user experience, engineers need to invest time to understand workloads and bandwidth requirements for all critical functions and build policies and thresholds that match those requirements.

As business go through rapid transformation, any planning you make at the start of the project will go through significant changes throughout the implementation phases resulting in significant cost and time overruns.

Engineers need to invest time to understand workloads and bandwidth requirements for all critical functions and build policies and thresholds that match those requirements.

Unrealistic expectations of cost savings and last mile availability myths

Availability of low cost, high speed connectivity options in all your business locations is usually a mixed bag. A common misconception is that there is a wide variety of low-cost, high-speed, competitive connectivity options at any given location.

One of the reasons SD-WAN is considered as an alternative to more expensive legacy networks is the fact that you can use multiple, low cost connections simultaneously, providing equal or better performance and reliability than the legacy network.

For wireline connectivity, the best case scenario is that you have a single cable carrier and a single ILEC (Incumbent Local Exchange Carrier) that can provide service at an address. For additional connectivity options you'll need to start moving into other metered or more expensive connections for additional capacity and redundancy.

When looking at TCO (Total Cost of Ownership), you may find that SD-WAN is more of a cost neutral upgrade than a cost savings exercise depending on what your bandwidth requirements are. That's not all. The cost and effort that goes into managing multiple bandwidth providers and holding them accountable for connectivity SLAs at every one of your locations is usually underestimated.

Interoperability with existing infrastructure can become a nightmare

Most SD-WAN implementations are OTT (Over The Top). This means that the SD-WAN technology is added to the existing network infrastructure and replaces a function or functions in the environment.

OTT deployments re-purpose components of legacy infrastructure including connectivity, data centers, firewalls, and switches. This requires the engineers designing the solution to anticipate any interoperability challenges that may arise when deploying SD-WAN.

For example, if the SD-WAN solution includes an existing MPLS underlay, chances are you need the MPLS (Multiprotocol Label Switching) provider to cooperate in the design and deployment of the solution. This may not be an easy task if you're deploying SD-WAN to ultimately displace the expensive carrier solution!

While all of these challenges can make SD-WAN implementation daunting, you can overcome these hurdles if you plan ahead and set realistic goals for the implementation.

We identified four practical design considerations that enterprises should focus on before choosing an SD-WAN solution and moving forward with the implementation.



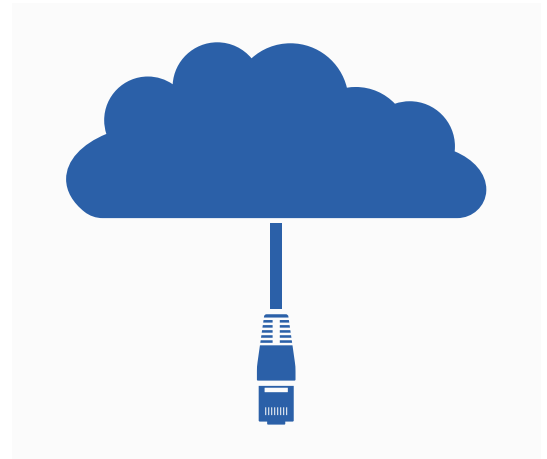
DESIGN CONSIDERATION

1 INITIAL DESIGN AND DEPLOYMENT

All SD-WAN platforms come bundled with orchestration software that allows network administrators to create policies that can be applied to network devices. This is a big step forward for companies that historically required a deep knowledge of their CLI (Command Line Interface) to configure and maintain the network configuration.

SD-WAN orchestration software, be it cloud-based or on-premise, follow a common theme:

- ✓ Create base configuration template(s) with administrative/device information, traffic flow preferences, and when applicable, security policies.
- ✓ Instead of configuring a device with a configuration script, a policy is applied, and device specific variables are configured/applied.



You may now begin to wonder, “That isn’t zero touch!”
You are correct!

While SD-WAN platforms have toned down the zero-touch marketing messaging, the reality is that while SD-WAN is potentially lower touch than legacy edge platforms, it does require initial setup and configuration of policies along with setting up interoperability with legacy network infrastructure.

The big problem that these configuration templates solve for is eliminating many of the pitfalls of manual policy administration including:

- Painstaking effort to build CLI (Command Line Interface) configuration syntax
- Demands deep platform expertise
- High margin of error
- Labor intensive to deploy
- Difficult to maintain ongoing configuration uniformity
 - Every time a change is made, staging templates must be adjusted
 - Changes will inevitably fail during the initial push for some devices, requiring an exception process to follow up and ensure changes are applied to ALL devices
 - Difficult to audit configuration compliance of devices

APPLYING NETWORK CONFIGURATIONS - OLD VS NEW WAY

Let's say you have a 500 location WAN, each with a router at a branch location that has the same template configuration. Your business is migrating from the standard Microsoft Office suite to Office 365. To support this effort, you need to update the configuration of your routers to ensure a smooth transition.

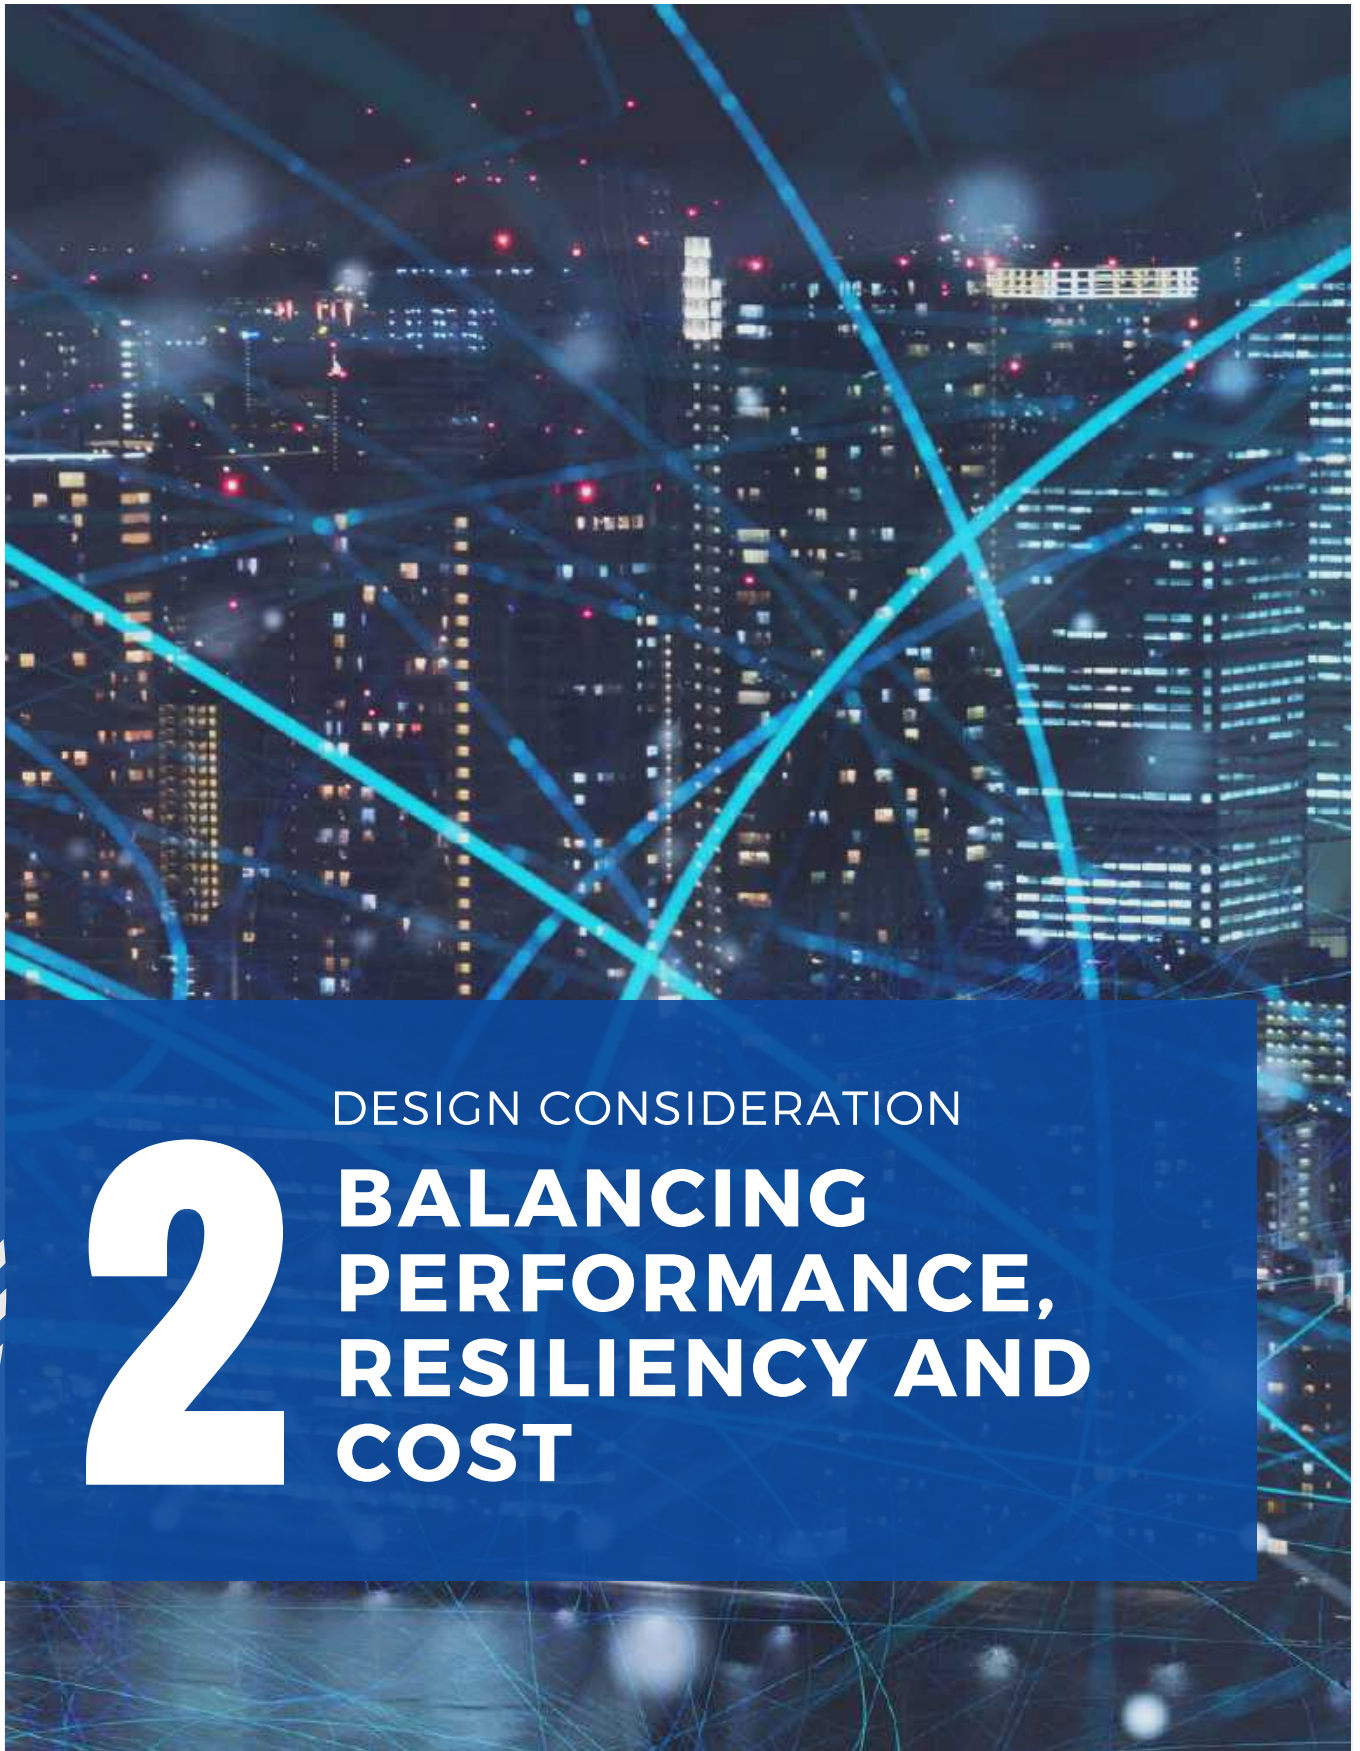
The old way (CLI)

- Identify necessary ports, protocols, and external hosts for Office 365
- Create command syntax to allow traffic from internal hosts to the new external hosts with the necessary ports and protocols
- Using an SSH (Secure Shell) scripting tool, run a batch job to write the change to the routers based on a current device inventory
- Run a batch job using the SSH scripting tool that includes show commands to ensure that the change was applied to every device
- Document failures
- Re-run the batch process until all devices are updated or perform manual updates
- Update any staging or configuration templates for deployment or support

The new way (Template)

- Create a new rule in the SD-WAN template you wish to change
- Select the Office 365 application
- Set the parameters for the traffic (internal trusted hosts and traffic policy)
- Save the template (Apply to devices by group or globally)

As you can see, the key change with a template-based model is in the reduced level of effort and margin for error. You will still need deep knowledge and expertise to enforce best practices for security and performance. The new tools just accelerate the time to implement changes and reduce end user frustration.



DESIGN CONSIDERATION

2

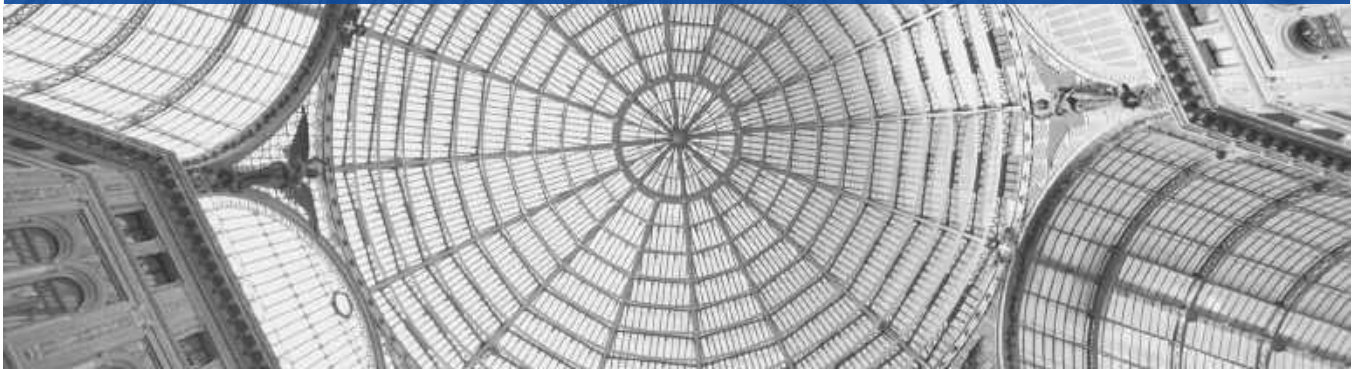
**BALANCING
PERFORMANCE,
RESILIENCY AND
COST**

One of the key drivers of SD-WAN adoption is the continued deployment of bandwidth intensive applications, real-time collaboration, and cloud adoption/SaaS.

The promise of SD-WAN is that businesses can leverage lower cost connectivity and achieve better reliability and performance than ELAN (Ethernet Local Area Network) or MPLS networks. While this is probably true for some locations, remote or difficult to reach locations may not have low-cost high-speed connectivity available.

While SD-WAN is able to overcome some of the performance challenges associated with sub-optimal connectivity, your SD-WAN network will still be bound by the laws of physics in terms of throughput and latency of the underlay network.

BUILDING YOUR CONNECTIVITY STRATEGY IS FOUNDATIONAL



When investigating SD-WAN options, the fundamental consideration should be the underlay network that the SD-WAN platform sits on top of. Typically, enterprises have one of the following connectivity scenarios based on the connectivity Cost:



Hybrid with MPLS/ELAN primary and Broadband Backup



DIA (Fiber/Metro Ethernet) Primary with Broadband Backup



DIA (Fiber/Metro Ethernet) Primary with Cellular Backup



Dual Broadband



Broadband + Cellular or Dual Cellular

Note: Cost will vary based on **footprint, carrier, cellular data plan size, and volume.**

In addition to connectivity type and budgetary cost, two key considerations are:



The footprint of your locations



The availability of connectivity options

For example, if you are in large office buildings, you may find that on-net DIA is readily available for primary connectivity. If you are in remote standalone locations, significant buildout costs may be required to deliver DIA. The same is true in the case of dual broadband availability. With the continued improvement of cellular networks, many customers find that its speed and reliability is better than a secondary wire-line option with the added benefit that a backhoe cannot knock out both connections!

Every industry has different connectivity preferences. A tried and true approach for retailers is broadband primary with cellular failover, and even dual cellular where wire-line broadband options are inadequate. This works well in this highly competitive vertical because it is inexpensive, generally reliable, and meets the application needs of a typical retail store.

For non-retail verticals where even 30 seconds of downtime means significant business disruption and loss of revenue, it may be worth the additional expense to have an underlay network consisting of DIA, broadband, and cellular at every location with hardware redundancy on SD-WAN and LAN equipment.

BALANCING PATH SELECTION, STEERING & SESSION SURVIVABILITY WITH USAGE AND EXPENSES



Steering is one of the fundamental elements of SD-WAN as it contributes to the promise of increased speed and reliability. Steering is the SD-WAN software's ability to determine the best path for network traffic based on predefined business policies and active health monitoring of available connections.

There are two scenarios when it comes to WAN path selection

SCENARIO 1 : NEW SESSION IS ESTABLISHED



The first scenario is when a new session is established. A request hits the SD-WAN appliance to establish a new session and the SD-WAN software selects the best path based on link performance and business policy priority.

Different platforms have different measurements for how a link is selected for a new session, the big questions here is:

Does your design require load balancing across multiple links, or should certain traffic be pinned to a specific connection unless it is down or severely degraded (I.E. traditional failover)?

For example, if your second connection is metered (think cellular), you may only want to reserve that connection for certain types of traffic when the primary connection is completely offline to reduce potential overage charges. Additionally, if your primary connection is a high speed Fiber or Cable connection and your second connection is low speed DSL, you may only want to send traffic over the low speed DSL if the performance of the primary circuit drops below a predefined threshold.



SCENARIO 2 : LINK DROPS/DEGRADES WHILE THE SESSION IS STILL LIVE



The second scenario is when a session is already established on a link and the link drops or becomes degraded while the session is still live (like a telnet session or VoIP call). Session based steering does not allow a previously established session to move to another link. If session-based steering is the only steering method supported, this scenario would result in a dropped session or poor application performance.

There are two methods for active steering (session survivability) that are employed on some SD-WAN platforms:



PACKET-BASED STEERING

In this scenario, the SD-WAN software can move a session to an alternate link without severing the session. The telnet session or VoIP call is able to move to the non-severed or degraded link without dropping.



DUPLICATION

In this scenario, the session stream traverses multiple links simultaneously and if a single link is severed or degraded, it will continue on the link that is still up and performing well.

It is important to note that both packet-based steering and duplication only work in a bookended fashion where the SD-WAN fabric is on both sides of the WAN edge. These solutions are typically more expensive than options that do not include this feature. Also be aware that while this design increases the real-time survivability of key applications, it inherently requires more hub and spoke routing of packets which potentially increases risk of a larger scale outage unless each hub is very well designed.

Link steering, both initial and active are big differentiators between platforms. The important item to note is that when it comes to network design, there is not a one-size-fits-all approach.

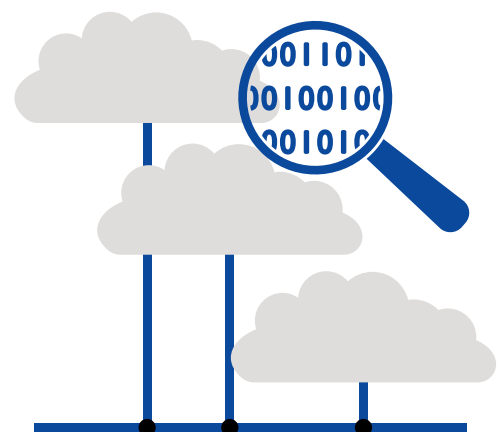
ASKING THE RIGHT QUESTIONS FOR SD-WAN COST-BENEFIT ANALYSIS



Both the underlay (connectivity), as well as the overlay (SD-WAN) vary greatly when it comes to cost and performance. IT leaders must look at the specific use cases that differentiate platforms and weigh the cost/benefit of each by asking yourself questions like:

- ❓ How often do link outages and brownouts occur on the connectivity I am using?
- ❓ Is active steering necessary based on the probability of link degradation mid-session, and what is the cost of that event?
- ❓ Is my second connection metered (cellular) and do I want to use that connection all the time?
- ❓ If my associate drops a phone call when the SD-WAN appliance re-routes traffic over an alternate connection, will the customer call back or is that lost revenue?
- ❓ Do I have highly compensated employees where the overhead expense associated with delayed file transfers is more expensive than a faster connection?

Once you have weighed your specific use cases, you can establish budget and strategy around connectivity. Only then are you ready to take an objective look at the dynamic traffic management benefits of different SD-WAN platforms, since now you are considering it through the lens of expense and usage management.




CASE STUDY

PRACTICAL APPLICATION OF BALANCING PERFORMANCE AND COST

One of our large retail clients routinely pulls video footage during normal business hours. Since this can be a bandwidth intensive exercise, downloading footage at some locations was causing high link utilization and poor overall network performance. While some locations have low cost connectivity with adequate bandwidth to download video footage without affecting performance of other critical applications, other locations only qualify for low speed broadband or expensive bonded T1 or DIA. Since the video footage needs to be available on demand during normal business hours, we needed to solve for the fact that:

- There were locations with insufficient bandwidth where the performance of critical applications was severely downgraded when the bandwidth-intensive app was running.
- There were budgetary constraints on last mile connectivity, eliminating the possibility of installing a higher capacity circuit



Our managed SD-WAN implementation allowed us to provision multiple low cost connections and configure business policies to load balance traffic across all of the connections, as well as traffic shaping policies to prioritize critical applications. The end result was sufficient bandwidth to provide a good end user experience for all applications without exceeding the allocated budget for network connectivity.



3

DESIGN CONSIDERATION

ONGOING VISIBILITY AND MAINTENANCE

SD-WAN orchestration software not only simplifies the initial configuration and deployment, it can be used to:



View real-time and historical application performance



Generate email or API alerts for outages or performance threshold breaches



Perform mass configuration changes based on templates

Different SD-WAN platforms have varying degrees of visibility and control, as well as some self-healing capabilities. While these are excellent tools for fine tuning network configuration and capacity planning, the following facts are still true of SD-WAN and legacy WAN technologies alike.

- Connectivity will go down
- Connectivity outages often require triage to repair
- Cellular connectivity is metered*

*Note: As of this writing, even “unlimited” plans have limitations on the types of allowed traffic and can throttle speeds over certain consumption levels

In addition, broadband and cellular connectivity typically come with best effort SLAs (Service Level Agreements). This means that even if you have a degraded link and can show the carrier evidence of a performance issue, there is often little that can or will be done to alleviate the problem.

One of the benefits of SD-WAN is that in a blackout or brownout scenario the software will attempt to:



Prioritize traffic so the most mission critical applications are given the most bandwidth



Send traffic over alternate links based on link health metrics



Use forward error correction to enhance data transmission reliability

The key consideration when evaluating the visibility and control capabilities of an SD-WAN platform is how you plan to use it.



If you have **made a significant investment in the speed and redundancy** of your underlay network, having visibility to enforce SLA and fine tune the load balancing characteristics of the SD-WAN platform may be a good use of time.



If you **have low cost links with cellular failover at branch locations**, it may not be important to routinely evaluate traffic flow policies to ensure you are getting the most out of your provisioned links.

All of these should be considered with the mindset of how you plan to consume the SD-WAN service as well. If you are considering a managed service, your measure of success will be network availability and performance, and proof of performance when making configuration changes and working trouble tickets.



If you are **considering DIY SD-WAN**, consider the time spent not only measuring and tuning the network, but also how you will be alerted of outages, your incident management process and tracking systems, as well as carrier and SD-WAN vendor escalation processes.



4

DESIGN CONSIDERATION

SECURITY AND SASE READINESS

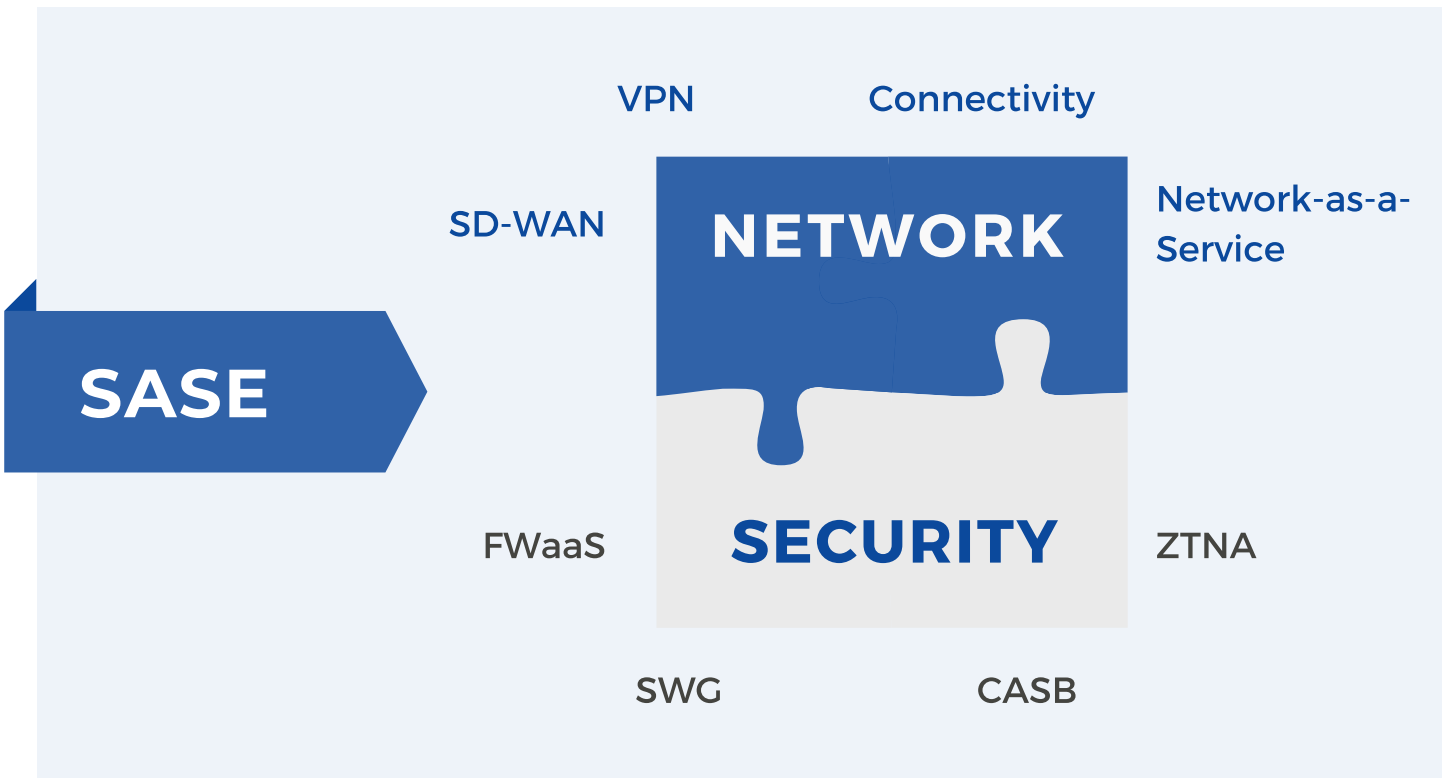
While the Gartner definition of SD-WAN includes service chaining for functions like WAN Optimization and Firewalls, the practical application of “service chaining” in most SD-WAN networks is routing traffic through a 3rd party firewall.

Early SD-WAN providers needed to solve for the fact that their built in stateful firewall did not meet the security needs of some organizations, especially since all traffic was no longer backhauled through corporate data centers for inspection. The initial options were to run a security VNF (Virtual Network Function) on the SD-WAN appliance or send traffic to a cloud-based security platform. Due to a low VNF adoption rate, most SD-WAN providers have focused on building integrations with cloud-based security providers. This not only alleviated the problem, it also helped give birth to the concept of SASE (Secure Access Service Edge).



While SASE is a best practice security framework that consists of SD-WAN as a core component, the market is rapidly evolving and will be largely left out of this brief.

What you need to know about SASE is that it is a framework that consists of a combination of network security functions like SWG (Secure Web Gateway), CASB (Cloud Access Security Broker), FWaaS (Firewall as a Service), and ZTNA (Zero Trust Network Access), alongside WAN capabilities like SD-WAN.



The main considerations for security are:

- ❓ Where do your applications reside?
 - Cloud providers like AWS or Azure?
 - Your own data centers?
 - 3rd party data centers like Equinix?
- ❓ Do you want security functions to be performed at your branch or in the cloud?
- ❓ Do you already have an investment in a security platform that you want to integrate?
- ❓ Do you want remote VPN users on the same security framework as your WAN?

CHOOSING BETWEEN SITE-LEVEL SECURITY VS CLOUD PERIMETER SECURITY

Many financial services organizations have thin clients and host applications on their own infrastructure. Since this architecture requires traffic to be sent to specific data centers that are customer owned, it makes more sense to enforce security at the customer hosted data centers rather than a cloud based firewall as it adheres to the “closed infrastructure” security approach.

For customers that do not have a hub and spoke design and all applications are SaaS or Internet, the decision becomes a matter of site level security vs. cloud perimeter security.

Organizations can realize financial and operational benefits from platforms that control SD-WAN and security on a single framework, administered on a single pane of glass. While a multi-vendor strategy for WAN and security is not inherently a bad choice, ensuring interoperability is proven and that you are able to administer access and permissions easily and effectively across the various control planes should be a key consideration.

BRINGING IT ALL TOGETHER: BEST PRACTICES FOR DESIGNING AND IMPLEMENTING AN SD-WAN SOLUTION



Implementing SD-WAN should be part of an overall strategy to optimize your network design. The table below lists some of the critical tasks in ensuring a smooth SD-WAN migration. This is not meant to be a comprehensive list but includes some best practices to illustrate the level of effort that goes into a successful implementation.

PHASE - NEEDS ANALYSIS

TASK	DESCRIPTION
Application Profiling Current State	Build a list of all applications on the network and where those applications currently reside.
Application Profiling Future State	Build a list of all applications on the network and where those applications will reside when the SD-WAN implementation is complete.
Application Requirements	Establish performance requirements for applications; required bandwidth, latency, and jitter thresholds.

TASK	DESCRIPTION
LAN Profiling	Evaluate segmentation and LAN requirements and determine if existing IP schema works in new environment.
Use Cases by Application	Determine use cases for SD-WAN functions (steering, session survivability, etc.).
Visibility Requirements	Determine what information is needed on network and application performance.
Administrative Requirements	What level of control is required and what expertise exists on your current team.

PHASE - SOLUTION DEVELOPMENT

Connectivity Profile(s)	Establish location type profiles and establish cost/performance thresholds.
Use Cases for Testing SD-WAN	If doing a PoC (Proof of Concept), use cases are critical in testing SD-WAN functionality, especially in a competitive situation.
Platform Rationalization	Based on the connectivity profiles, application requirements, and visibility/administrative requirements, choose platform(s) to be evaluated and/or deployed.
Configuration Templates	Based on the connectivity profile and application requirements, develop configuration templates for each location type.

TASK	DESCRIPTION
Transition Interoperability	Determine if and how migrated locations will communicate with legacy locations and services until the transition is complete.
Segmentation Design	Build LAN segmentation based on best practices.
Establish a Lab	Configure and deploy at least one device in a controlled environment.
Deploy Headend Equipment First	Data centers should be deployed first, this allows two way communication for real-world testing scenarios.
Security Configuration	Establish security configuration and ensure it aligns with business requirements and best practices (I.E. PCI, HIPAA, to name a few.).

PHASE - SOLUTION HARDENING

Install New Connectivity	If migrating to a new connectivity profile (I.E. from MPLS to Internet), testing in an environment that closely simulates the final product is critical. Things as simple as two circuits with different latency characteristics can throw off the performance of the SD-WAN platform.
Test Use Cases	Run through testing scenarios to ensure performance of the SD-WAN functions produces the desired result.

TASK	DESCRIPTION
Security Audit	Internal or external resources should evaluate the security posture of the newly designed solution.
Template Adjustment	As needed, adjust the configuration templates until use cases and security requirements are satisfied.

PHASE - SOLUTION DEPLOYMENT

Installation Documentation	Determine the process to update a branch location (physical, logical) as well as any head end adjustments to avoid problems like asymmetrical routing.
Connectivity Deployment	If new connectivity is in scope for the SD-WAN deployment, it is ideal for circuits to be installed before the SD-WAN device implementation.
Configuration Tuning	It is likely that you will uncover scenarios in the field that you were not able to simulate in the lab. This will require you to quickly pivot and adjust configuration templates as needed.

PHASE - ONGOING MAINTENANCE

Configuration Changes

Deploy changes like whitelist entries or new business policies for apps as required by your business.

Break Fix

Identify and resolve connectivity and device issues as they occur.

OS Maintenance

Review release notes for new OS revisions of the SD-WAN software and perform risk analysis based on known bugs, deploy OS updates based on risk analysis.





SD-WAN IMPLEMENTATIONS DONE THE RIGHT WAY

When you work with Interface, our consultative approach to network design and expertise in rolling-out performance-driven networks will give you a significant competitive advantage.

- ✓ **CUSTOMIZED DESIGN AND IMPLEMENTATION**
- ✓ **INDUSTRY-LEADING SLAS AND GOVERNANCE PROGRAM**
- ✓ **CARRIER AND TRANSPORT AGNOSTIC CONNECTIVITY**
- ✓ **CONSOLIDATED FLAT-RATE BILLING**
- ✓ **24X7 SOC AND NOC MONITORING**

INTERFACE IS
TRUSTED BY

GENESCO

SALLY BEAUTY

Michael's
MAKE CREATIVITY HAPPEN



James Avery
JEWELRY

LONG JOHN
SILVER'S



TruckPro

rubio's IHOP
COASTAL GRILL

AND MANY MORE...

GET A FREE CONSULTATION