ZIVARO

BRILLIANT IT

Its Time Your WAN Strategy Caught Up to Your Cloud Strategy

Aaron Shelley
Program Architect

BUZZWORDS EVERYWHERE

In a world constantly moving toward a cloud-first, or at the very least, cloud-heavy strategy, the connective tissue of an organization's technology footprint is more important than ever. Even as the advent and adoption of cloud solutions drives change at a breakneck pace, most organizations are still relying on legacy WAN infrastructure and design that has changed very little over the last decade. More and more organizations are realizing that they are "not in the data center business" and are choosing to focus on what is truly core to their business model. As a result, more and more organizations are consolidating their data center footprints into a few purpose-built facilities, rather than spread across regional locations. Couple this with the meteoric growth of cloud and X as a Service (XaaS) solutions, and connectivity issues have gone from painful to revenue impacting for most organizations. This in turn has resulted in many connectivity issues for locations with bandwidth and latency constraints being the typical culprits. Historically, this has meant either provisioning additional MPLS connections or increasing the capacity of existing ones. Both paths tended to be costly and slow with very little in the way of visibility or control of the part of the IT staff. Modern networks require a better solution. As the traditional application model has shifted to a predominantly cloud and XaaS model, the traditional WAN architecture (MPLS connection to a data center then on to the application location) is no longer the best suited strategy given the new traffic patterns driven by cloud solutions. It is also no longer simply about bandwidth requirements. Thanks to the evolution of XaaS and communication applications (UC, VOIP, etc.) the nature of the network requirements at all locations have changed, but the traditional MPLS-centric solution offers little-to-zero control over, or visibility into traffic prioritization.

Enter SD-WAN. According to a report by Global Market Insights in 2020, the global market size for SD-WAN is set to grow to nearly $30B by 2026[1]. In 2018 the total market for SD-WAN solutions was hovering around $1.5B so there is some impressive growth expected as cloud migration continues. From a quick glance, SD-WAN offers cost savings on the circuit front, increased control, flexibility, security, and agility for the WAN as a whole. All of this is absolutely true, however, just like architectural benefits of cloud, SD-WAN is not a STRATEGY, it is part of one. To realize the benefits of SD-WAN, all too many organizations are falling prey to the inherent problems of buzzword compliant architectures which frequently introduce more instabilities and costs than they eliminate. SD-WAN is a simple enough concept, but it is typically not simple to implement. For most organizations there are simply too many things that the network touches (aka dependencies) for it to be simple. Just as cloud-centric architectures must be approached with a different mindset than traditional on-premise architecture,

1. https://www.gminsights.com/pressrelease/software-defined-wide-area-network-sd-wan-market

SD-WAN must be approached differently than a traditional WAN. The most common approaches to implementing SD-WAN are to either simply "bolt" it on top of the existing MPLS, adding redundant paths via DIA or cellular, or to completely rip out MPLS and run on purely SD-WAN with DIA connections. Either approach is going to result in issues. Here are the issues that we see most often.

"I Was Told There Would-Be Savings…" – your CFO… probably

This happens most often with the bolt-on approach. You are likely to end up wondering where your savings are. There is a potential for some savings in a bolt-on approach, but only if at least one MPLS connection is replaced with a DIA connection. Even then, the cost of the SD-WAN appliances as well as their implementation must be offset. SD-WAN can absolutely result in more savings, including on the hardware front, but ONLY if there is proper planning

Somebody Needs a Nap

Why is the network team so stressed and overworked? This is another issue that most often accompanies a bolt-on approach. A new "brain" has been introduced into the network to specifically make decisions as to what traffic takes what path and with what priority. Combine this with the inevitable unknowns that lurk in the network, and you have a recipe for new surprises. If proper planning has not considered the current solution as well as the business needs, things are going to be messy. While the concept of SD-WAN is very simple, its affects upon the network will not be.

Can You Hear Me Now?

As communications solutions has advanced, the latency sensitivity of the associated traffic has grown. While a lack of QoS can be overcome with enough excess bandwidth, that approach does not work when the internet is involved. Many organizations have discovered this immediately after implementing SD-WAN (with both common approaches). Sometimes it is a QoS issue, sometimes it is a previously unfelt misconfiguration in the network (such as a synchronous routing). Communication solutions are also the main reason that most organizations will decide to have a WAN architecture including both SD-WAN & MPLS.

To Mesh or Not to Mesh

To Mesh or Not to Mesh?

This pitfall is reserved for organizations with a larger number of locations. Typically, SD-WAN is going to be implemented with at least two WAN connections. This means that if an organization has 50 locations with two connections each and wants to fully mesh the WAN, then all SD-WAN appliances will need to support at least 98 tunnels; with 100 locations, the requirement becomes 198. The math is (Total Sites–1)*(WAN Connections per site). This can often result in organizations choosing to size the SD-WAN appliance for a location by the tunnel requirement rather than the throughput. However, a Hub(s) & Spoke architecture should not be ignored. While fully meshed has many benefits, it can frequently have financial constraints.

SO, WHAT DO WE DO?

These will seem like no brainers. However, more often than not failed or painful SD-WAN implementations are caused by a miss on one or more of the following: have a defined business outcome, have a plan, and do not go it alone.

Something that we see IT teams miss in general all too often is to ensure that major changes being made can be directly pointed to an intended business need or improvement. If an organization is looking at implementing SD-WAN, they should have a defined business-impacting reason like measurable cost savings or improved application availability. There are plenty to choose from that may be applicable, but some measurable outcome(s) need to be decided. Remember not to fall prey to simply having a technical reason or goal. Yes, SD-WAN is a technical solution, but it needs to answer a business need. Tech for the sake of tech can keep the attention of the staff implementing the technology, but it is also the fastest path to merely being a cost center and to the perception that something is being done because it is the latest shiny object that can cure a deficiency. Remember that the goal is not the plan. Implementing an SD-WAN solution is likely to affect almost every technical aspect of your organization, and in doing so proper planning is even more vital. Given all of this, do not go it alone as there is no plausible reason or need to do so. Whether you seek advise from a trusted peer, OEM vendor or partner, you should speak with someone for outside viewpoints. SD-WAN is no longer bleeding-edge, but as we hope we have made clear, it is still not mature to the point of being a basic utility. There is someone in your network that has worked with and likely successfully implemented SD-WAN who can help provide new and different perspective on what a successful transformation would look like.

WHAT TO CONSIDER

With all the above said, here are the things that we would advise any organization to consider prior to implementing an SD-WAN solution:

- **Current Contracts & ETLs** – Are any going to block the ideal solution? This is easy to overlook but can be a complete block to moving forward (e.g., if an organization just renewed a multiyear contract on redundant MPLS connections at all sites).

- **How Many Locations** – This will likely be the main driver in whether to implement a fully meshed or hub and spoke architecture. Remember that if you have 50 sites with dual connections, then you would need to be able to support a minimum of 98 tunnels per site. Most organizations of any appreciable location count are going to be utilizing a hub and spoke architecture. And with the shift to Cloud and XaaS, often at least some of the hubs are in 3rd party locations (Colocation DCs, IaaS Solutions, Clouds, etc.)

- **Standard Classifications of Locations** – This simple planning step often greatly simplifies the design, planning and deployment of an SD-WAN. Classifications should include standardized needs for uptime and capacity. These two needs will then drive how many WAN

connections and what size and type are appropriate for the location. That analysis will then make it clear what SD-WAN appliance is appropriate as well (do not forget to consider the tunnel counts).

- **Typical Traffic Types** – This and ETLs are the most likely drivers as whether or not MPLS will continue to be part of the WAN architecture. The answer you arrive at may very well not be an all or nothing answer. Often organizations have locations with differing enough needs that not all warrant the more expensive connection.

- **How to Implement** – Aside from the decision of bolt-on or rip-and-replace, the timing aspects of rolling out the changes to sites must be considered. This can often be driven by timing of existing contracts on WAN connections as well. There is not a single correct vs. incorrect path here, it is case by case based on the unique attributes of the environment.

EXAMPLE

The following is an actual real-world example of a customer that made the decision to add SD-WAN to their technical infrastructure. Given the frankness of the information provided, the customer name will not be provided, and some points will be kept purposefully vague.

Customer Profile

- Regional Organization with B2B & B2C Operations
- Heavily Regulated Industry
- Over 250 Locations Spanning Both Metropolitan & Rural Areas
- Exponential Growth via Acquisition (Over 4X in 3 Years)

Customer Technology Profile – pre SD-WAN Implementation

- MPLS-only WAN Architecture (~50% of sites with single connection)
- Hub and Spoke WAN with a single external facing ingress/egress site for the overall network
- Cisco-based Network Infrastructure
- 3 On-Premise Data Centers, 2 Third-Party Data Centers & 1 Public Cloud Provider
- No VDI In Use
- Heavy SaaS Use
- Just Delving into Public Cloud Utilization

Customer Business Goal/Need – that lead to the decision to implement SD-WAN

Executive and board directive to establish the organization as a true “Next Gen” provider in their industry while growing the organization into one of the 50 largest in the country through frequent and sizable acquisition.

Customer Known Technology Pain Points

- Frequent Throughput Capacity Issues at Ingress/Egress Site
- High Rate of Acquisition of Sites with Disparate Technical Stacks
- High Lead Time on MPLS Changes
- Short Window to Convert Acquired Assets – typically 2-3 days for complete integration of systems
- High Cost on MPLS Circuits
- Lack of Redundant WAN Connections at Many Sites – lack of ability to obtain in rural areas
- Higher than Average WAN Latency – due to single hub model

Customer Plan to Achieve the Goal

Establish a nimbler overall technical architecture utilizing SD-WAN to eliminate many issues which had to date, greatly slowed conversions of acquired assets as well as slowing or at times preventing the implementation of industry leading solutions and functions.
Other changes were made as well, such as the implementation of VDI and building of a B2B public cloud solution, but the focus of this the SD-WAN.

Customer Requirements for New Architecture

- Reduce Overall IT Spend Over the Next 3 Years
- All Sites Have 2 WAN Connections Minimum – diverse paths required
- All Sites Maintain an MPLS connection – End to End QoS on primary connections determined a MUST
- No Single Point of Failure
- No Site Isolating Event Can Disrupt Other Sites
- Move to True Active / Active Data Center Architecture – original was Hot / Very Cold in functionality
- Leverage Cellular WAN Connections as Needed – for rural locations with limited ISP options
- Single Solution Viable in Azure, IaaS, On-Premise Data Centers & Remote Office Locations
- Public Internet Traffic Direct Egress/Ingress from All Locations with Acceptable Security
- Manageable with Current IT Staff – no additional head count

Customer Chosen SD-WAN Solution

Cisco Meraki was selected after significant research and comparison between several leading SD-WAN solutions. The ease of implementation, flexibility and integration with existing systems were the main driving factors. Cisco's Meraki product line is one of the most popular platforms that providers (like Zivaro) are offering in managed SD-WAN service bundles. Given the dearth of options out there, that is saying a lot as many of the tier-1 telecom providers seem to be selecting Meraki as (one of) their managed SD-WAN offering(s). In 2019, Cisco's SD-WAN portfolio including the Meraki platform was the overall winner in CRN Magazine's evaluation of software-defined networking solutions.[2] Meraki has some very elegant design elements that make it relatively simple to manage and it has terrific analytics through its Meraki Insight tool for how the network is performing. You can find information on this important piece of the solution here.

https://meraki.cisco.com/products/meraki-insight/
https://documentation.meraki.com/MI

2. https://www.crn.com/slide-shows/mobility/crn-s-2019-products-of-the-year/20

Unexpected Complications During Implementation

Note that all the below complications existed within the original architecture prior to SD-WAN being introduced. These complications were not "felt" until the SD-WAN overlay effectively "ripped off the band aid." SD-WAN caused none of the issues, but the initial perception was "SD-WAN broke our network," but that was about as accurate as saying that no longer taking Tylenol caused someone that was already sick to get a fever. Neither caused the problem, they simply removed the effect that was masking the symptoms.

- Spiking Latency – due to asynchronous routes
- Decline in Call Quality – due to QoS misconfigurations & asynchronous routes
- Security Configuration Gaps – unrelated to SD-WAN, found during implementation
- Marked Increase in Bandwidth Utilization in Major Sites – due to asynchronous routes

Outcome of Implementation

Throughout this very lengthy project, Zivaro was able to fully imbed and integrate with the customer to operate as a part of their internal team at every level from executive leadership to architecture, and engineering to deployment in the field. The ability of both organizations to trust each other was key to the success of this project and has continued to be so as Zivaro has assisted in the conversion and integration of acquired locations, as well as overlaying the monitoring of the existing internal operations center with Zivaro's Managed Services NOC. This partnership allowed the customer's internal staff to focus on strategic initiatives rather than up/down response efforts and predictive modeling (not to mention the ability for them to effectively eliminate their need to be on call).

The final implementation consisted of five (5) hubs and 260 remote office locations with all hubs fully meshed to one another and all remote office locations connecting to all hubs. There is a total of five (5) standardized location classifications. Once Zivaro was able to guide the customer through the unexpected complications listed above, the architecture was able to reach a new normal. After one year of operation the new WAN made up of MPLS & Meraki SD-WAN in combination has resulted in fewer network outages (on a per site and overall basis both), lower latency and a distinct improvement in call quality.

CONCLUSION

Application usage has evolved over time and recently it has done so at an ever-increasing rate thanks to the advent and adoption of software definition (a.k.a., the cloud). Systems have become more and more distributed, and workforce users also have become increasingly distributed. The network however, and the WAN in particular have not evolved at the same pace, at least from an adoption standpoint. Treating the evolution of the network as importantly as the march to the cloud is critical for successful arrival at a transformed architecture. The options and abilities exist with great solutions like Cisco's Meraki SD-WAN platform. If your organization is not taking steps to adapt and evolve to meet the requirements of modern users and applications and doing so at all levels of the tech stack, then you can be certain that many of your competitors likely are. While every organization's situation may not be evolve-or-die, it is at the least a state of evolve or face the consequences of declining business capability without proper attention to the network's necessary and parallel evolution.