TechVision RESEARCH

The Path to SD-WAN

Abstract

As organizations move to become Digital Enterprises, we see the traditional network, storage and compute infrastructure is increasingly software defined. The requirements for speed, flexibility, scalability and the movement from capital to operational expenses are driving this trend. Innovations such cloud computing, virtual machines, containers, edge-computing, and virtual storage are removing the limits on the compute/storage by providing independence from the underlying physical assets. Software Defined Wide Area Networks are driving the transition of networking (transport) in much the same way.

Software Defined Wide Area Networks (SD-WANs) promise to help enterprise networks run faster, better and at a lower price point while improving security. These purported benefits have led to rapid growth in the SD-WAN space with CAGR in the 70% range. Cloud, mobility, IoT, and edge computing are disrupting traditional networking models for connecting users and devices to services, applications, and data. This movement to the edge also supports the high SD-WAN growth.

Historically SD-WANs represent the 4th generation of Wide Area Networks and totally transform the way WANs are designed and built. Most research to date in this area focuses on the benefits of SD-WANs, but we'll not only cover the value of SD-WANs, but the reality of how and when SD-WANs be used most effectively, architectural approaches and associated trade-offs. Every enterprise network is unique and selecting and implementing an SD-WAN strategy is a complex process.

SD-WANs provide improved network performance and security as applications can be hosted anywhere in the cloud and users are mobile and everywhere. The software part means that routing and security can be run on commodity hardware, virtualized infrastructure, or within a cloud hosted environment. The future is clearly SD-WANs and this report will walk you through the benefits, architecture, roadblocks, planning process, migration and deployment guidelines.

Author:

Sorrel Slaymaker
Principal Consulting Analyst
sorell@techvisionresearch.com

Table of Contents

Executive Summary

SD-WANs represent the 4th generation of WANs and are a total transformation in the way WANs are designed, built, and managed. IP networks are more important now than ever to ensure performance and the security of applications that are hosted anywhere (on-premise, in the cloud) and mobile users who are everywhere.

The challenge in building and managing SD-WANs is that every enterprise and government agency has unique requirements and approaches for building their networks. With 5x network bandwidth growth projections, a requirement to ensure 99.999% reliability, and support for the 24/7 Digital Enterprise, WAN transformation is a significant undertaking.

The promise of SD-WANs is to make networks faster, better, cheaper, *and* more secure. While there are some trade-offs in these variables, early adopters of SD-WANs report they see benefits in all areas. This early success is driving the substantial expected growth in this space. The table below summarizes the differences between current traditional WANs and the future-state moving to SD-WANs.

Architecture	Today's WAN	Tomorrow's SD-WAN
Security Model	Perimeter and Trusted Inside	Zero Trust, w/o Borders
Routing & SLAs	Static & Based on Link	Dynamic & Based on Apps
Provisioning	Manual	Automated & Zero Touch
Service Provider	Few Selected Providers	Any Provider
Capacity Allocation	Purchased in Advance	Elastic, On-Demand
Transport	MPLS	MPLS, EPL, Internet, 4/5G
Management	On-Premise	Cloud
Hardware	Proprietary	Commoditized, Virtual
Demands for Growth	Branch to Data Center	Cloud, Mobile, IoT, Edge
Design	Hairpin Through Data Center	Direct User to Application
Reliability	99.9%	99.999%
Application Visibility	Probes & 3rd Party Tools	Built Into SD-WAN Router

Table 1: Differences between WANs and SD-WANs

While SD-WANs are a significant improvement over today's WANs, there are drawbacks. The most significant challenge is that there isn't an SD-WAN protocol standard. This results in all current solutions effectively being proprietary, which in-turn, results in vendor lock-in. There are also over 60 vendors in the market selling SD-WAN solutions, and Cisco, the traditional market leader, is lagging in this market thus opening the door for enterprises to evaluate alternative vendors.

To design an SD-WAN, there are many architectural considerations. First, the business architecture defines the requirements and ownership of the solution, in which many enterprises are choosing a managed service. Second, the technical architecture is critical in designing the network underlay and the new SD-WAN overlay. Third, the security architecture involves moving towards a Zero Trust model that involves a more granular segmentation strategy and a decision as to where to provide the Internet offload. Fourth, a Unified Communication (UC) architecture is important since video consumes the majority of the network bandwidth and voice is one of the most critical and network sensitive applications. Finally, open source and IPv6 are future considerations for an enterprise SD-WAN strategy.

Figure 1: Benefits of SD-WANs

Moving to SD-WAN is a significant network transformation. Many enterprises will try to bolt on SD-WAN to their existing Wide Area Network and that will limit their ability to take advantage of this new technology. The equivalent analogy is designing a new electric car from the ground up like Tesla versus bolting on an electric motor and battery to an existing car. We recommend the Tesla approach in this context.

Introduction

Wide Area Networks (WANs) connect users and devices across many locations to services, applications, and data that reside in data centers and cloud providers. Every 10 - 15 years, WANs go through a revolution on how they are designed and built. Early WANs used modems to transmit data over phone networks. Second generations used frame relay and ATM, and third generation WANs used MPLS. Software Defined WANs (SD-WANs) are the fourth and latest generation; they are transport agnostic, meaning they can use MPLS, Internet, Ethernet Private Line (EPL), and 4/5G cellular.

Since the enterprise WAN is usually the bottleneck in delivering high performance, low latency applications, a lot of attention is given to reducing the bottleneck. Figure 2. shows users and devices on the left side and the services, applications, and data they consume on the right side with a WAN in the middle.

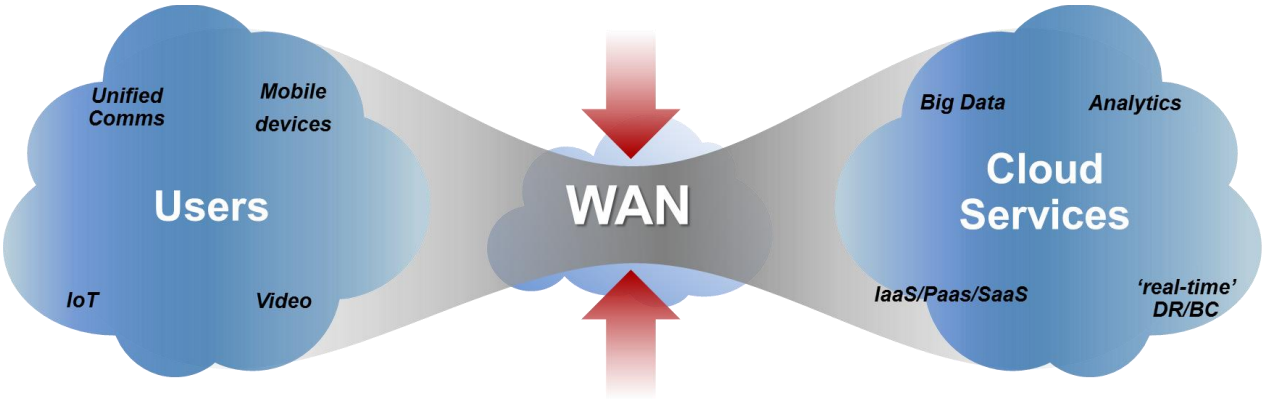

Figure 2: WAN as the Bottleneck Between Users and the Applications They Consume

Software is "eating the world" and has made its way to networking with Software Defined Networking (SDNs). Software Defined Wide Area Networks (SD-WANs) are a subset of SDNs and are getting market attention because WANs are a major bottleneck for enterprises and consume a large portion of the networking budget.

Enterprise networks are a conglomeration of many different networks including data centers, branch offices, and campuses, each with their Local Area Network (LAN) and connected to the Internet, Cloud Service Providers, VPNs, and Partner networks. WANs are used by enterprises to create a private network that inter-connects all these locations and networks together as shown below in Figure 3.

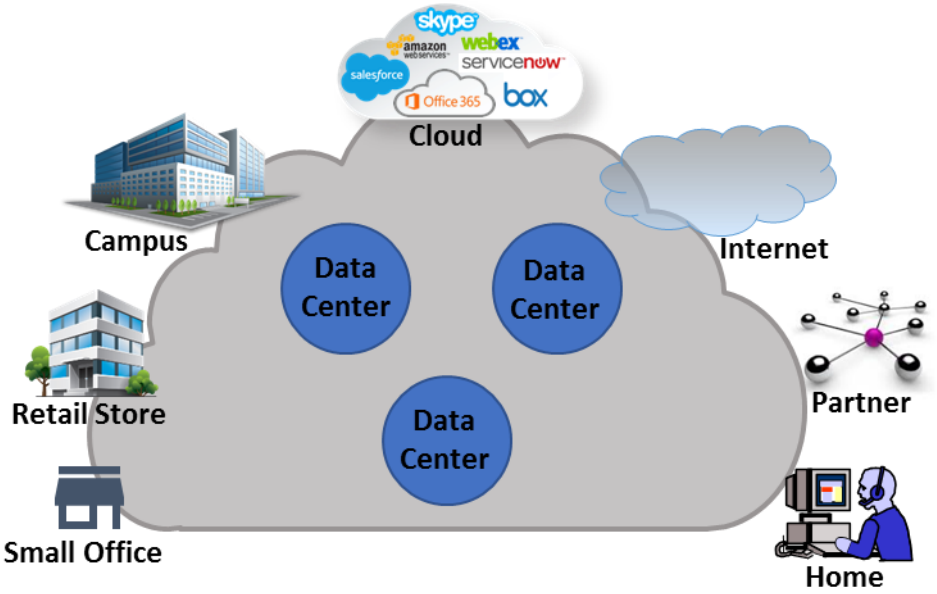

Figure 3: A WAN Inter-Connecting Different Locations & Networks Together

The enterprise WAN is typically 10% of the IT budget. The primary components of a WAN are:

1) **Access** – The last mile connectivity between a site and a network service provider. Traditionally this access was copper (T1, DSL, or DS-3), but in the last decade many sites have been connected with fiber. Cellular LTE and forthcoming 5G wireless connections are becoming more common. Access can be shared between many businesses in an area such as Internet DSL and Cable, or dedicated access. Access accounts for **60%** of WAN costs and also are the most problematic as the local copper or fiber connections often get severed during construction or other projects. The old saying of a "fiber seeking backhoe" applies here, creating interrupted services for a day or more.
2) **Transport** – The network service provider's private network such as MPLS, EPL, 4/5G cellular, and dedicated Internet. Transport represents about **15%** of WAN costs, and many enterprises report that the cost for private MPLS and dedicated Internet from the same service provider is the same. Many reports claim the savings of going to SD-WANs is in moving away from using MPLS. This is not the case: most of the savings are in going from dedicated access to shared access, such as a T1 MPLS to Internet DSL.
3) **Edge Router** – The hardware, software, and enhanced functions such as firewall security and/or WAN optimization to provide the network "steering" and quality of service guarantees. Also, the management and troubleshooting tools are critical since the network is guilty of slowing an application down unless it is proven otherwise. This represents

about **15%** of WAN costs. These costs are high in part because the enterprise routing market has been dominated by Cisco for two decades.

4) **People** – Architects, engineers, technicians, and project managers who design, build, and support a WAN are the final **10%** of WAN costs. Depending on the business requirements, the sophistication of tools and automation, and the rate of change, enterprises find that they need one person for every 50 - 100 branch offices or 10 - 20 inter-datacenter and cloud connections.

Traditionally, the scope of the WAN was relatively straight forward since services, applications, and data were all housed in the company's data centers, all employees worked within one of the company's office sites and they used devices that were company owned and managed. But this is all changing with users being mobile and using their own devices while services, applications, and data are moving to the cloud. Furthermore, instead of just relying on private and expensive MPLS connections, enterprises have the options of using 4/5G wireless, EPL, and Internet with the options of dedicated or shared access.

The network latency required for the next generation of applications will shrink down to five milliseconds.

Low-latency application requirements will push MPLS and traditional network WAN architecture into the graveyard. In a hyper-converged digital world, milliseconds matter. People using virtual reality and self-driving cars require near instantaneous data, computing, and connectivity. The network latency required for the next generation of applications will shrink down to five milliseconds.

Cloud 2.0 architectures will force a totally new network and security architecture. In today's enterprise networks, backhauling traffic to a central data center is common. Enterprises centrally deploy and manage VPNs, mobile device management, proxies, and private/public internetworking. They backhaul traffic since putting all the network security controls on the edge is too expensive. With highly distributed data centers in Cloud 2.0, security will also become more distributed.

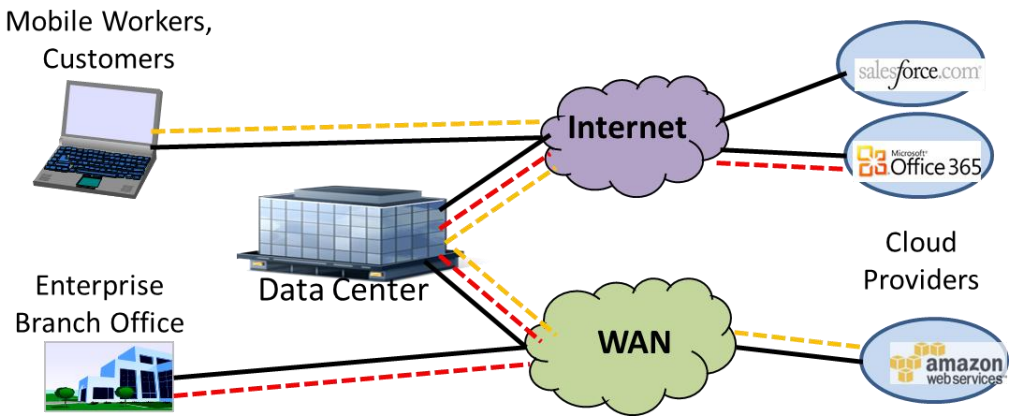

Figure 4: Example of Backhauling in an Enterprise Network

The cost of putting networking and security software at the edge of networks is dramatically falling as startups bring new approaches to the market. Instead of a network architecture of edge, distribution, and core — where all the routing intelligence and security controls are at the distribution layer — future network architectures will have an intelligent edge running on any commodity hardware/VM with no distribution layer. The core in this type of architecture is a superfast IP forwarding plane.

Today's SD-WAN implementations perpetuate backhauling versus connecting users directly to the services they're consuming. SD-WANs use pre-established (static) tunnels that add a 30% bandwidth overhead. In many cases, one of the tunnels goes to a cloud security stack which still forces backhauling, which is the bane of network architecture when every millisecond counts.

> *Future network architectures will have an intelligent edge running on any commodity hardware/VM with no distribution layer.*

The reason users experience slower cloud application performance in the office than they do at home is because of all the backhauling that occurs in the corporate network today. In one example, a contractor found downloading a 100-megabyte file from Office 365 SharePoint took two minutes and 10 seconds in the office but only 13 seconds at home.

From his office in Seattle, the contractor would log on to the VPN to gain access to the guest Wi-Fi network, which had an IPsec tunnel to the Wi-Fi controller. The tunnel ran over an SD-WAN, through a data center 1,000 miles away, and then back out to the hosted O365 site in Seattle. The network latency at home was 8ms to O365, while in the office it was 52ms for small packets and 77ms for large packets. The large packets were fragmented and also delayed by the encryption process of running through three different IPsec tunnels (VPN, Wi-Fi, and SD-WAN) even though the contractor had a secure TLS connection to O365.

Also, in Cloud 2.0, expect cloud providers to get into the networking market so they can ensure the performance and security of their applications all the way to the user. As every millisecond counts, cloud providers will rely on over-the-top (OTT) delivery using Internet transport to connect with their users and build private internetworks for their internal use.

Why WANs Are Important, Unique, Growing, Yet Slow to Change

WANs will continue to play an important role for all enterprises. While the Internet, a network of over 1,000 networks, interconnects everyone and everything, a private WAN is still required to ensure end-to-end network performance and security. Large cloud providers such as Google, Facebook, Amazon, and Microsoft have huge private WANs that are bigger than the large Network Service Providers (NSPs) such as AT&T, BT, and Tata. They use their private WAN to ensure the best performance, cost, and security for their services.

The Network is More Than a Utility

Many CIOs consider their IT infrastructure, including networking, to be the equivalent of a public utility. The Infrastructure as a Service (IaaS) market leaders Amazon, Microsoft, and maybe Oracle have furthered this notion with their on-demand, consumption-based offerings of compute and storage.

Similar to Maslow's hierarchy of needs at the physiological or basic needs level, the utility aspect of networking is just the foundation. Large enterprises and IaaS cloud providers that treat their networks as a utility tend to dramatically overbuild their network. They will then put their own layer 3 network, usually tunnel-based, on top to interconnect their layer 2 networks. These networks lack the end-to-end intelligence, security, and distribution that an all-layer 3 network that is session state aware, can provide. The network vision is to enable applications to dynamically and intelligently allocate bandwidth as required.

Maslow's hierarchy is a layered approach from the basic of human needs to self-actualization and achieving one's full potential. This same approach is analogous to the needs of an enterprise and its networks.

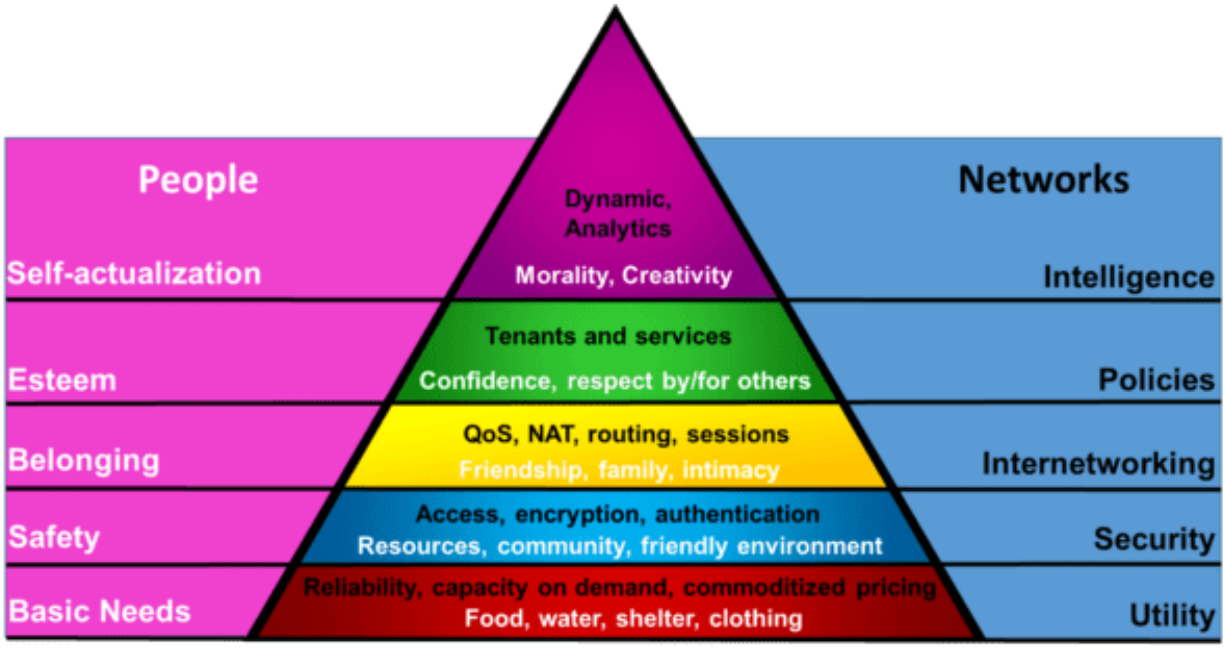

Figure 5: Maslow's Hierarchy of Needs/the Needs of an Enterprise and its Networks.

Attributes of networks that make them more than a utility are:

- **Bidirectional** – Being able to consume network traffic in both directions
- **Any-to-Any** – Ability to talk directly to any other device on the network
- **Multiple Delivery Options** – Wireline and wireless across many paths, including the last mile
- **Intelligence** – Providing QoS, rate limiting, and dynamic multi-path routing
- **Security Controls** – Complex security controls including access, authentication, encryption, segmentation, and logging with a Zero Trust architecture
- **Application Aware** – Understanding and optimizing the network for the applications that run on it

The network first must have all the attributes of a utility such as high reliability, capacity on demand, and commoditized pricing. Then, the network must have the self-awareness and intelligence to dynamically allocate the network to user and application demands.

A great application is defined by the user experience, and the only way to guarantee a great user experience is to ensure that the network is intelligent, not just a utility.

A great application is defined by the user experience, and the only way to guarantee a great user experience is to ensure that the network is intelligent, not just a utility

Why Every Enterprise WAN is Unique

The future looks bright for network architects. Every enterprise WAN is unique, and this will continue, even as networking technology moves to software. It is estimated that Cisco and its resellers get over $50 billion in professional services revenue every year due to the customization and complexity of every enterprise network. The global managed network services market is around $130 billion with a CAGR of 6 - 7%

There is no such thing as a simple turnkey SD-WAN solution

Every enterprise network also comes with its own requirements and priorities. This reality means there is no such thing as a simple turnkey SD-WAN solution. In fact, WAN networking will get even more complex for several reasons:

- A disappearing network perimeter with users, devices, services, applications, and data everywhere
- Very low latency requirements to support virtual reality and other next generation, real-time applications
- The need to be always available 24/7/365
- Heavy bandwidth requirements driven by video that is always on including surveillance, virtual offices, remote monitoring in IoT, and augmented reality.

Through SDNs, better orchestration tools, and market innovation, enterprise WANs will hopefully get simpler, but unlike compute and storage, each enterprise WAN will always be unique. Attributes that influence the architecture of every enterprise WAN include:

1) **Geographical Footprint** – The location of offices, partners, manufacturing, and the list goes on. Some businesses are concentrated in a small geographical area such as a hospital system that may have eight hospitals and 90 clinics in one county. Others are distributed all around the world. Local bandwidth is cheap, resulting in an enterprise network that tends to overbuild their network versus global bandwidth that is expensive and has high latency. The result is WANs that are tightly managed and utilize WAN optimization technologies.

2) **Business Vertical** – Enterprises with numerous small branch offices such as banks, insurance or commodity retail versus those with a few large sites such as manufacturing, pharmaceutical, higher education, etc. Some enterprises are a mix of both. Small branch offices utilize traditional copper-based access (T1/E1, DSL, Cable) while large sites have optical access at 1/100th of the cost per Mbps. Organizations with a few large sites typically deploy a layer 1/2 WAN utilizing technologies such as VPLS, while the enterprise with thousands of small sites utilize layer 3 technologies such as MPLS or Internet VPN to scale.

3) **Business Applications** – Some businesses are transaction-oriented, so their network bandwidth needs are simple. Other businesses are interaction-oriented and require video collaboration and sharing of very large CAD, MRI, or other large files, which consumes 100x the bandwidth. The location of the applications — in the cloud, private data center, or within the office — impacts the network design.
4) **Security Policies** – The security position of being minimally compliant, industry standard or best in class has a direct correlation on network security requirements. Network access control, encrypting all data in motion, and network perimeter strategy (a few Internet exit points and trust the internal network versus zero trust networking) all impact the amount of network routing and efficiency of the network. Too many networks backhaul traffic when going from the enterprise network to the public Internet. Also, different industries have regulatory requirements that influence the security required, such as credit card transactions running on their own segmented network.
5) **Lifecycle Management** – Most enterprises have contracts with network service providers, along with network technology vendors, and they replace each contract in different six- to eight-year timeframes. Each major refresh and change calls for a significant investment in time and money, and each major change utilizes the technologies and products available at the time. A major network refresh takes three to four years. After such a refresh, an organization wants to harvest the investment for another three to four years.
6) **Organization Dynamics** – The decision makers and priorities of the decision makers are different. Some businesses centralize all IT infrastructures into a single team to try and get economies of scale, while others have IT controlled by each major business division. With centralized IT, some businesses charge back directly to each business division, while others use a shared pool business model. Different decision makers value cost, reliability, and performance differently which shapes the overall network architecture.
7) **Culture** – IT architecture philosophy including:
    a. Best of breed versus single vendor
    b. A defined architecture versus project-by-project planning
    c. Treating the WAN as a utility versus business differentiating
    d. Insource versus outsource some or all network functions and/or management

There are surely other factors that also contribute to making every enterprise WAN architecture unique. Networks are complex, and it is incumbent on every good network architect to keep things as simple as possible while also not trying to follow any preset design. IoT, Cloud, Mobile, and digital business are sure to keep network architects busy for the foreseeable future.

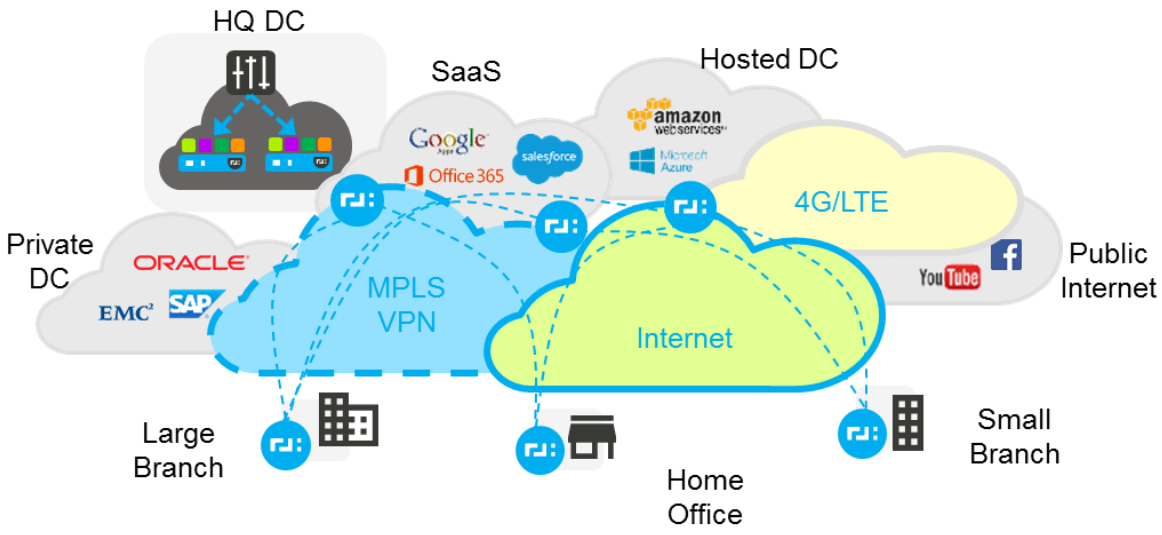

Figure 6: Example WAN that Inter-Connects Different Locations & Cloud Services

In 2019, eight large U.S. retailers, encompassing over 2,000 stores, embarked on new SD-WANs. Though each had thousands of branch offices in the U.S., their business objectives for moving to SD-WAN were vastly different, which resulted in vastly different architecture and solutions for each. Each used a different SD-WAN vendor, some chose to use LTE, and some chose to have local Internet breakout. Some relied on a third party for implementation and support while others did it internally. In the architecture section, all the options will be reviewed. The point is that going forward WANs will continue to be unique to every enterprise. In fact, with software, the ability to customize grows, enabling more uniqueness.

Network Utilization Will Grow 5x by 2024

Network architects need to budget and design their networks for peak busy hours of utilization, which will grow 5x in the next few years. Peak utilization is growing significantly faster than average utilization, driven by our real-time world that reacts to major events.

Major events can be supply chain disruptions, failure of an assembly line, major weather system, a global pandemic, a product defect reported by multiple customers online, etc. When a major event occurs, people need to see what is going on and collaborate with others to manage the situation. It is at these critical moments that the network must have the capacity to handle real-time and near real-time video, which will consume 95% of network bandwidth.

One example of bandwidth growth is for contact center agents. Currently, enterprise network architects plan on utilizing an average of 200Kbps per agent with 140Kbps for applications delivered via a virtual desktop with 720dpi screen resolution and 60Kbps for Opus-based high-quality voice. The busy hour is 2.5x busier than average call volumes. By 2024, as contact center agents utilize augmented reality and video-based customer interaction at a resolution of 1080dpi, the average bandwidth per agent is predicted to be 2Mbps. The busy hour will be 3.5x busier

than average, driven by social media and near real-time events such as a sales promotion of a product/service or a defect.

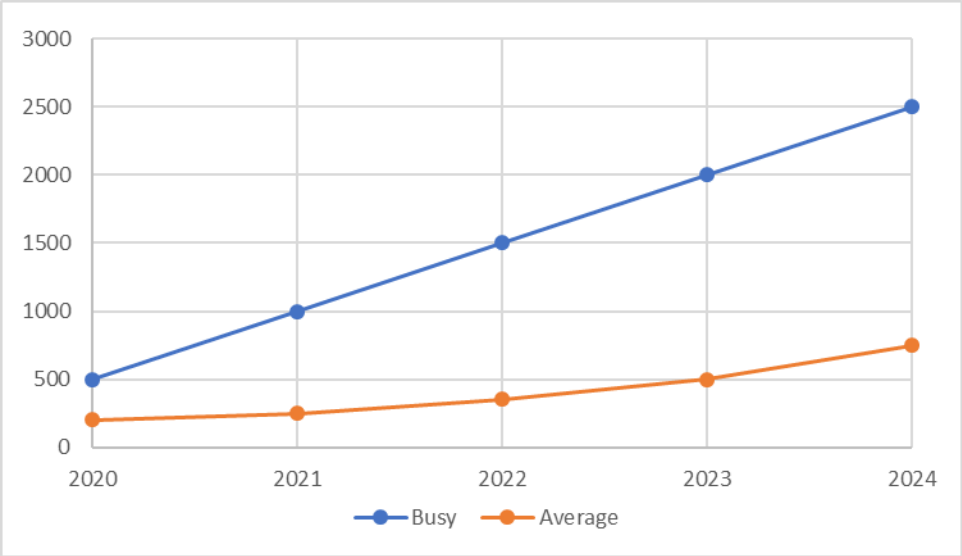

Figure 7: Average and Busy Hour Bandwidth in MBPS per Contact Center Agent

According to Cisco, busy hour Internet traffic is forecasted to grow 5x by 2024, compared with average utilization which will grow 3x in this time period. 2020 is also the timeframe where 5G cellular will become widely deployed and IoT is forecasted to grow from 6.5B to 30B devices. Network architects should adopt the following strategies to deal with this explosive growth:

Busy hour Internet traffic is forecasted to grow 5x by 2024, compared with average utilization which will grow 3x in this time period.

1. **Overbuild network access** – For critical business sites and sites with over 50 people, ensure Ethernet access over fiber with the ability to easily grow from 10Mbps to 10Gbps. Broadband network access works only for smaller and non-critical sites.
2. **On demand pricing** – Change telecom contracts from buying a peak amount of bandwidth to unlimited and bill based on 95th percentile of utilization, which is the model used with inter-networking between service providers.
3. **Do not rely on a single network** – Instead of bringing in one to two networks, enterprises should connect to carrier neutral, co-location services. They should also take advantage of the 1,000+ global fiber networks and buy bandwidth at wholesale rates.
4. **Realize that CAC & QoS have limitations** – Dynamic and intelligent management of traffic when congestion occurs instead of random early packet drops. Not all critical

voice, video or application sessions are equal; some require a segmentation strategy based on security trust and business criticality, not just DSCP classification.

5. **Improve reporting & analytics** – Instead of just monitoring utilization, tracking TCP & UDP packet drops, which is a better indicator of application performance than tracking bits per second.

Too many network architects base bandwidth forecasting on average utilization. They then rely on call admission control and quality of service to manage traffic during the peak business hour. While this methodology worked in yesterday's transaction-based networks, future networks are interaction-based, which consumes 5x more bandwidth.

Why Network Refreshes Take 3 - 4 Years

Today's networks are very complex, and enterprises have been conservative in making any major changes, for fear of causing an outage that disrupts the business. Most network projects are tactical in nature to meet short-term needs such as adding another cloud provider and putting in firewalls in the internal network to improve security. This leads to further network complexity, higher costs, and a fragile environment that can break easily. Every so often, the network needs to be redesigned from the ground up.

CIOs and business leaders are busy and will only decide to refresh the network if they are experiencing a lot of pain such as outages, poor application performance, or network hacks/attacks. Once a decision is made to refresh the network, the process starts.

This is the standard process that enterprises go through to upgrade their network:

1) **Get Resources** – 3 - 6 months – The existing network team still needs to manage the day-to-day needs of the business and to fix ongoing outages while supporting new applications. Hiring a team can take upwards of a year since these resources are scarce in the market and bringing in contractors or managed service partners means that a lot of the knowledge will walk out the door at some point.

2) **Define requirements** – 2 - 4 months - Interview business leaders, IT architects, network managers, and IT operations for what the network requirements are currently and in the future. While this can appear at a high level to be fairly easy, every company has different requirements and different priorities, which makes this step very important, time consuming, and feeds into the next step.

3) **Issue RFPs** – 6 - 12 months – Issue RFPs to carriers and network equipment providers to get a sense of the right services, equipment, and professional services. This process can be as much political as technical. Select 1-2 network service providers and 1 - 4 network equipment providers (routers, firewalls, IPS, WAN optimization).

4) **Create an Architecture** – 3 - 6 months– Architecture creates a blueprint that all stake holders can understand for a design that has competing priorities. Cost, reliability,

performance, security, and operational simplicity are competing priorities, and gaining consensus takes time for people to understand the trade-offs and come to consensus.

5) **Pilot** – 3 - 4 months – Test and pilot services and products to ensure they deliver as advertised. This is especially important for new technologies such as SD-WAN. Besides making it work, ensure the administration and operational troubleshooting and reporting tools are in place. Many organizations do this step last, not first, and experience a lot of pain in their migration.

6) **Implementation** – 9 - 18 months – This always takes longer than planned because:

    a. **Change** – The business requirements, leadership, and markets change on a regular basis. Enterprises are notoriously bad at multi-year projects with priorities, re-organizations, and budgets changing each year.

    b. **Outages** – Impacts to business from outages can grind a project to a halt and require the review of everything – architecture, vendors, staff, and operations model.

    c. **Resources** – Because of the complexity and hardware centric nature of networks, they require a high touch implementation with both central and on-site resources. Coordinating resources on project plans that can fluctuate is a nightmare and most upgrades require 50% more resources than were planned.

7) **Operations** – 3 - 6 months – Getting the day-to-day processes, reporting, and change management with the new infrastructure and handing it off the existing IT operations team to take the support calls.

8) **Retiring the old network** – 3 - 6 months – While this may sound obvious, fully retiring something within enterprise IT is difficult. There is always one feature or application that is better on the legacy environment due to years of tuning. Plus, taking out the old equipment and finding things that were not identified in the discovery and requirements process.

Going forward, a 3 to 4-year implementation is not acceptable. Networks, like businesses need to be elastic and scale up and down as required.

Going forward, a 3 to 4-year implementation is not acceptable. Networks, like businesses need to be elastic and scale up and down as required. One of the great things about moving networks to a software model is being able to separate them from the underlying hardware. As more network routing and security need to be done, the associated virtual hardware can scale as required.

The Promise of SD-WANs

SD-WANs are not a technology looking for a business problem, but business problems needing a better way of interconnecting sites and networks. These sites can be office buildings, private homes, the over 1,000 cloud providers, partner locations, and/or retail stores. There are a tsunami of business drivers coming, along with a number of networking inhibitors that are forcing organizations to turn to a more revolutionary transformation of their WAN versus the evolution that has occurred to date.

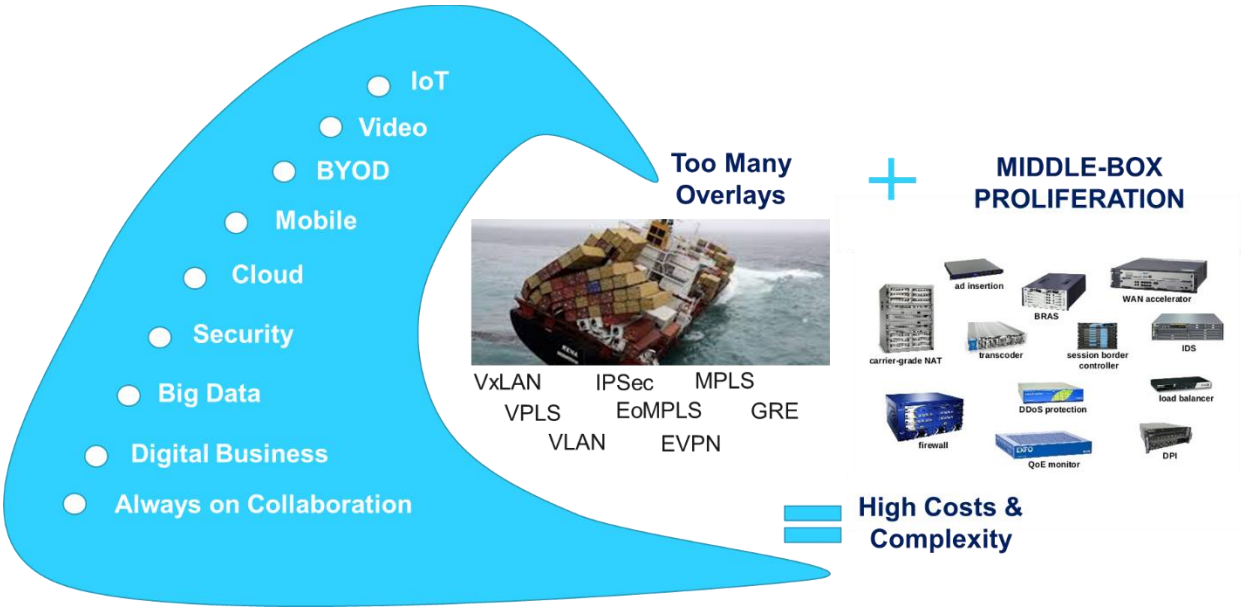

Figure 8: Tsunami Converging on Legacy WANs

Figure 8 shows on the left, the tsunami of requirements from IoT to Always on Collaboration that are driving additional WAN performance and security requirements. On the right, all the tactical, point protocols, overlays, and associated middle boxes that make today's WANs complex and costly.

Why SD-WANs?

The promise of SD-WANs is to be hardware, transport, service provider, and software SD-WAN vendor agnostic. SD-WANs provide improved network performance and security as applications can be hosted anywhere in the cloud and users are mobile and everywhere. The software part means

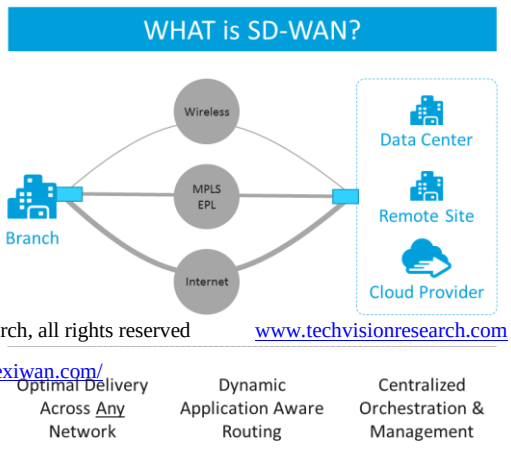

that routing and security can be run on commodity hardware, virtualized infrastructure, or within a cloud hosted environment.

SD-WANs are transforming networks and empowering enterprises to make their WANs:

1) **Faster**
    a. **Greater network performance** by directly routing users and devices to services, applications, and data without backhauling and the associated latency. Being transport agnostic allows enterprises to buy bandwidth cheaper and remove the bottleneck of the traditional WAN. SD-WANs are application aware and can provide sub-second rerouting to a better path during a brown or blackout of one of the network links.
    b. **Agile** – Near real-time changes and the elasticity to grow and shrink as workloads and associated business requirements change in this digital world. Gone are the days of fixed contracts and designs that inhibited flexibility.
2) **Better**
    a. **Automated** – Simple with zero touch provisioning and more automation with GUI-based management systems and doing away with the CLI.
    b. **Integrated** – DevOps tools so that networking bandwidth and associated routing and security policies can be provisioned dynamically as compute and storage change.
    c. **Highly Reliable** – Moving from 99.95% to 99.999% reliability where the network downtime to a site moves from 4.38 hours a year to 5.26 minutes. 24/7/365 digital businesses need to be working at all times and that failover from one network to another now needs to be less than 2 seconds.
3) **Cheaper**
    a. **Lower operational expenses** – Being service provider and transport agnostic, which means that instead of just using 1 - 2 NSPs and MPLS only, an enterprise can choose one of thousands of NSPs that has the best network at the best price in a given region, as well as mix and match MPLS, Internet, EPL, 4/5G Cellular. Lower maintenance costs and fewer human resources with automation also provide some operational savings.
    b. **Lower capital expenses** – Hardware agnostic and running on commodity hardware, virtual machines, and/or in the cloud.
4) **More Secure**
    a. **Zero trust architecture** – Providing 1:1 micro-segmentation between users and devices communicating with services, applications, and data. Whitelist security policies of what is allowed in a least privileged model instead of a blacklist model

that defines where network traffic cannot go. Anomaly detection provides early detection if a user is misbehaving or a device has been infected with malware.

b. **Encrypting data in motion** – Ensuring that all data on the WAN is encrypted, whether it is at the application layer using TLS or at the network layer using IPsec or other network encryption algorithm.

If enterprise IT does not provide an elastic, agile, cost-effective network that inter-connects everything dynamically, then business leaders will look for alternative options.

Digital enterprises are utilizing cloud, mobility, IoT, and edge computing to disrupt their markets by driving innovation and speed. This in turn is driving enterprises to rethink how to design, build, and operate their WANs that connect their users (customers, partners, and employees) and associated devices with services, applications, and data which reside everywhere in the cloud and need to be available at all times.

Architecture	Today's WAN	Tomorrow's SD-WAN
Security Model	Perimeter and Trusted Inside	Zero Trust, w/o Borders
Routing Policy & SLAs	Static & Based on Link	Dynamic & Based on Apps
Provisioning	Manual	Automated & Zero Touch
Service Provider	Few Selected Providers	Any Provider
Capacity Allocation	Purchased in Advanced	Elastic, On-Demand
Transport	MPLS	MPLS, EPL, Internet, 4/5G
Management	On-Premise	Cloud
Hardware	Proprietary	Commoditized, Virtual
Demands for Growth	Branch to Data Center	Cloud, Mobile, IoT, Edge
Design	Hairpin Through Data Center	Direct User to Application
Reliability	99.9%	99.999%
Application Visibility	Probes & 3rd Party Tools	Built Into the SD-WAN Router

Table 2: Comparison of Today's WAN versus Tomorrow's SD-WAN Architecture

Many business leaders express frustration with how slow, rigid, and expensive the enterprise network is compared to what they experience on their home network. With the rise of shadow IT, networking is the last thing that enterprise IT controls end-to-end. But if enterprise IT does not provide an elastic, agile, cost-effective network that inter-connects everything dynamically, then business leaders will look for alternative options, such as having the cloud provider deliver both their service and the network.

Elasticity is one of the advantages of the cloud and something that business leaders also expect in their network. A common network elastic billing method is 95th percentile based - Burstable billing. Attributes of network elasticity are:

- **Ability to scale up and down within seconds** automatically based on demand/load
- **Paying for what is consumed/used/needed** without any long-term contracts/commitments
- **Hyper-connected** with the fiber and cross-connects that are already in place at hundreds of global co-locations to thousands of service providers (network & cloud)

Application-Centric Networking to Make Your Network Faster and More Reliable

Every critical application on an enterprise network needs additional intelligent and dynamic network services to optimize performance and ensure security. In the enterprise, best-effort networking isn't good enough. An enterprise has to add services on top of its common underlying IP network to support the specific requirements of critical applications. For performance and security, these requirements include:

1) **Performance**
    - **Quality** – Especially for external voice along with video and web conferencing
    - **Reliability** – 99.999% uptime in this digital 24/7 business world
    - **Measurements** – Being able to monitor and track performance
2) **Security**
    - **Regulations** – Being in compliance with HIPAA, PCI, GDPR, and others
    - **Protecting Data** – Consumer, financial, private, confidential, regulated
    - **Control** – HR and business rules on what employees and application can do

Today's IP networks segment local traffic into VLANs and then over the WAN, assign specific QoS classes and VRFs to optimize TCP/UDP flows, and in the data center use VxLAN for segmentation and rely on overbuilding for QoS. A few examples of today's segmented networks are:

1. Voice – Real-time communication with its own network
2. Telepresence – Immersive room-based video systems
3. Payments – Credit-card authorizations
4. Guest WiFi – Non-employee network access

The problem with today's hodge-podge of technologies is that it doesn't allow an end-to-end view, which is becoming especially critical since mobile users and cloud-based applications are on different networks. To solve this problem, networks must become more intelligent. Intelligence comes by moving up the stack to Layer 5, the session layer, where intelligent services reside (see diagram below for an overview of the OSI model and the services provided in the session layer).

The session layer, the most critical layer for an intelligent Application Centric Network (ACN), provides the glue between applications and lower-level network functions.

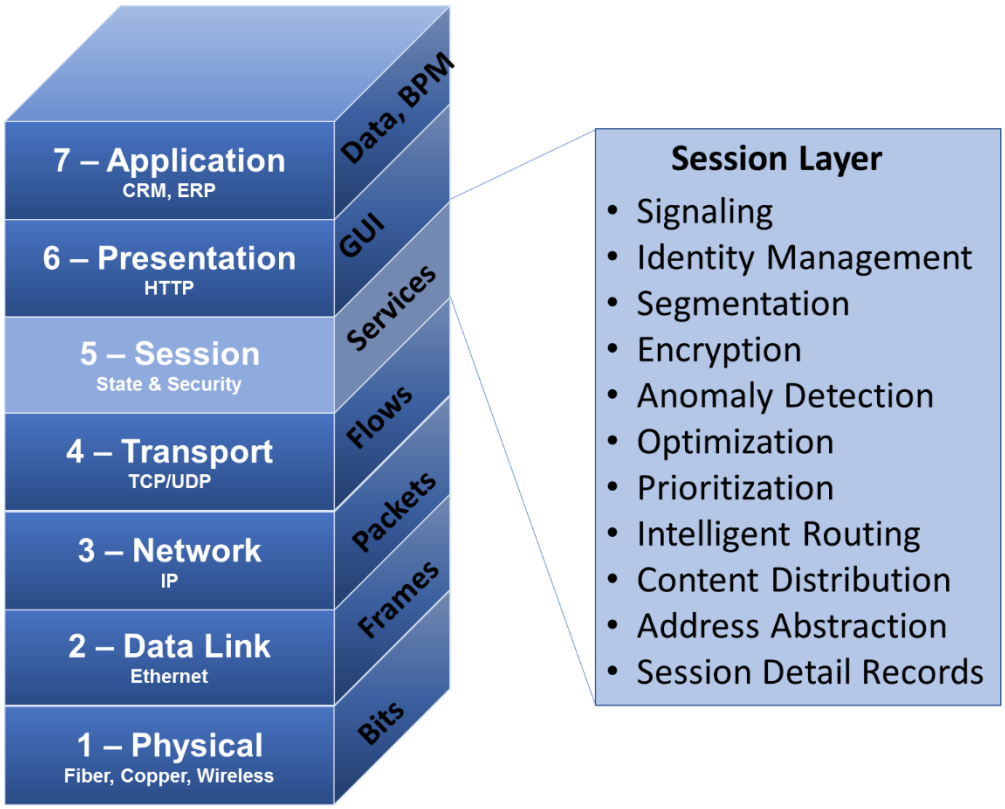

Figure 9: OSI Reference Model

The Session layer provides the mechanism for opening, closing, and managing a session between end users and applications. Sessions are stateful and end-to-end, which provides more granular network and security controls for application services. Firewalls, proxies, session border controllers, WAN optimization devices, load balancers, and caching/content delivery networks all manage the network state and provide higher-level networking and security functions. Instead of requiring the need to bolt on all these "middleboxes," network routers must provide these functions natively in next-generation networks.

A sample of services at the session layer fall into the same buckets of:

1) **Security**
    a. **Signaling** – Ability to request and secure network resources at the start of a session, which is a requirement for zero-trust networking

b. **Identity Management** – Integrating with directories to be able to verify users and devices

c. **Segmentation** – The rules on who is allowed to access what, and keeping the attack surface to a minimum

d. **Encryption** – End-to-end TLS and key management

e. **Anomaly Detection** – Understanding network use with alerts on misuse or malware

2) Performance

a. **Optimization** – Managing TCP/UDP flows such as window sizing and rate limiting

b. **Prioritization** – Dynamic network prioritization and controls, not only at the start of the session but throughout the entire session as other sessions come and go

c. **Intelligent Routing** – Using multiple paths with mid-session stateful failover

d. **Content Distribution** – Multi-cast support to distribute videos and files and supporting content at the edge of networks

e. **Address Abstraction** – Mapping the naming schema used in applications and directories to network IPv4/6 addresses and TCP/UDP port numbers

f. **Session Detail Records** – Monitoring and managing how network resources are consumed and accounting for network usage for planning and billing purposes

A lot of the focus of the SD-WAN market is to provide the glue between applications and the network. The challenge is that SD-WANs use tunnels and overlays such as IPsec and VXLAN, which are Layer 2 and Layer 3 based. These overlays don't work well through firewall/NAT boundaries, so they lack end-to-end user-to-application performance and security controls. SD-WAN vendors also partner with the session players to provide many of the stateful, intelligent network services, but finding one that provides all services is difficult.

Some will argue that intelligent routing is a Layer 3 function such as using BGP. While routing is a Layer 3 function, intelligent routing has state and end-to-end performance as well as security controls. For instance, if a link has a burst of errors or high jitter, a Layer 3 routing protocol will not re-route the application unless the link goes down. Layer 3 routers prioritize BGP keepalive packets that monitor a link health at the highest level, and thus are not impacted by network congestion, just network outages.

Just like next-generation firewalls are moving further up the stack, next-generation networking is doing the same. The ACN provides all the intelligent session layer features each critical application requires to ensure performance and security requirements are met.

Intent Based Networking to Make Your Network Better

One of the many (oh so many) "holy grails" for IP networking is the ability for an application to automatically request network resources and security. Intent-based networking (IBN) is the latest iteration, following past failures such as RSVP, IGMP, and NSIS. We'll now describe what Intent-based networking is, how is benefits the enterprise and how it benefits the organization and barriers to success.

Intent-Based Networking Defined

1) **Translation and Validation** – The system takes a higher-level business policy (what) as input from end users and converts it to the necessary network configuration (how).
2) **Automated Implementation** – The system configures the appropriate network changes (how) across existing network infrastructure via network automation and/or orchestration.
3) **Awareness of Network State** – The system ingests real-time network status for systems under its administrative control and is protocol and transport-agnostic.
4) **Assurance and Dynamic Optimization/Remediation** – The system continuously validates (in real time) that the original business intent of the system is being met and can take corrective actions (such as blocking traffic, modifying network capacity, or notifying) when desired intent is not met.

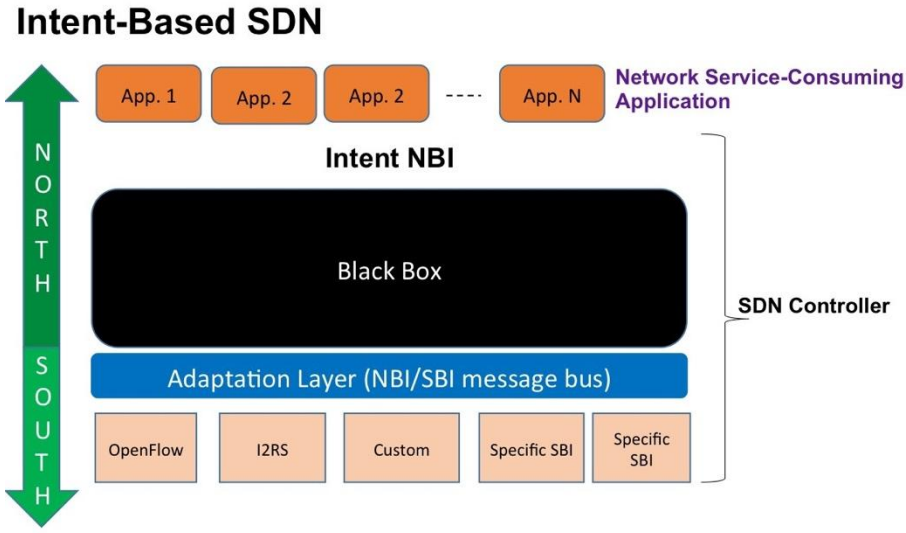

Figure 10: Intent Based Networking Framework

A few reasons why IBN will not work unless major network changes are made:

1) **Network Address Translation (NAT) Boundaries** – Users and the applications they consume are on different networks a majority of the time. These networks are separated by firewalls and NAT boundaries, such as IPv4 and IPv6, and lose their end-to-end

session-state awareness and controls, along with the ability to identify the traffic based on the initial source IP address.

2) **Non-Granular Differentiated Services Code Points** – A packet header value that is used for the Quality of Service requested for traffic, such as high priority or best effort delivery. The challenge with this is that a lot of the traffic on the network looks the same being either encrypted with TLS or video traffic. For instance, a room-based video between executives can have the same DHCP marking of AF41 just like a desktop video session between two coworkers.

3) **Lost in Translation Between Network Numbers and Application Words** – Ethernet and IP addressing are based on numbers and application use words for their programming logic. Translating between the two is difficult, with the added difficulties of NAT listed above and data center networks where north/south traffic may be IP address-based and east/west traffic is Ethernet address-based.

4) **Access Control List (ACL) Hell** – The underlying network is still controlled by many devices (routers, firewalls, load balancers, WAN optimizers) with each using ACLs to control network performance and security, hop-by-hop. Because a given ACL only applies to devices with a contiguous address space, the same ACL may be used on thousands of routers to control the same service across many disparate networks. IBN masks the underlying network complexity with orchestration tools, but it does not solve the fundamental IP networking problem.

5) **Signaling** – In order for an application to request network resources and security, a signaling mechanism needs to be used prior to the application using the network. While Cisco recently announced their plans on IBN, they do not have a signaling mechanism in their routers and firewalls.

Many SD-WAN vendors are using flows and session to help address the above limitations to automation and simplification in today's WANs. Some of the key enablers include:

1) **Session-Orientation** – Unlike a network flow, a network session has a unique session ID that enables it to pass routing and security policies through NAT boundaries and provide end-to-end network control and monitoring. Session Detailed Records provide the industry's best network performance and security analytics.

2) **Named Services** – Define all applications, services, and users on the network using words based on tenants and services. This naming scheme allows for granular network controls that are hierarchal and can be easily read by humans and machines. Address-independent routing provides location-independent services and user mobility.

3) **Single Software Stack** – Access all network functions (routing, firewall, load balancing, WAN optimization) in a single software stack that has Netconf and REST APIs that are easily controlled by SDN orchestration tools such as Ansible and Puppet.

SD-WAN platforms are relying on automation and a GUI to make networking simpler. The goal is to get away from the complex Command Line Interface (CLI). Some vendors have been more successful at doing this than others.

Reduce WAN Costs by 50% To Make Your Network Cheaper

Just like SIP trunking cut voice transport costs by 50%, SD-WAN can do the same, which is why there's so much hype. There have been enough business case proposals and implementations of SD-WAN to be able to put a stake in the ground that SD-WANs can cut WAN costs in half. Typically, WAN costs are about 10% of an enterprise's overall IT budget. Anytime a project can save money, improve performance, increase reliability, *and* provide additional security controls, it would seem that it would be a slam dunk.

> SD-WANs can cut WAN costs in half.

But just like with SIP trunking 10 years ago, there are many who are still reluctant to believe the magnitude of opportunity with SD-WAN. If you talk to sourcing managers, some of them think that issuing an RFP for traditional WAN services will save 25% and that implementation costs to move to SD-WAN do not outweigh the risks and effort. If you talk with network engineers, they tend to see the SD-WAN opportunity only for small sites where broadband Internet can deliver enough bandwidth and replace T1/MPLS. Many do not see large sites that require greater than 500Mbps as appropriate targets for SD-WAN.

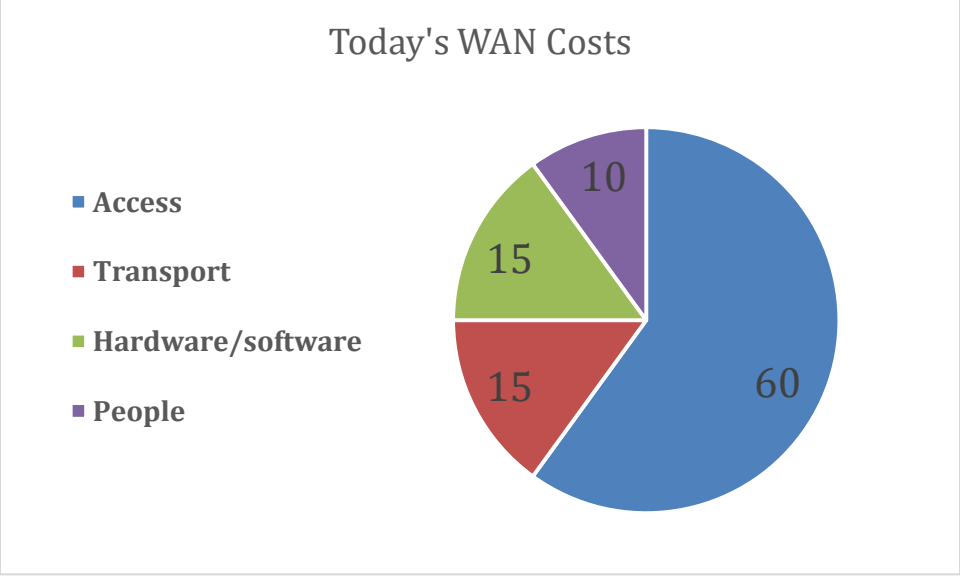

Figure 11: Today's WAN Cost Breakdown

These are the primary, pragmatic steps U.S. enterprises can take to achieve savings in the 50% range. TechVision's Top 4 Cost Saving Principles are listed and described in rank order as follows:

1) **Buy Wholesale Access and Be NSP Independent** – No one network service provider has complete coverage. Access, whether fiber or copper, is 60% of WAN costs. Network, fiber, and Internet service providers sell access to others at wholesale rates that are significantly less than heavily discounted retail rates. Earthlink (now Windstream), Global Capacity (now GTT), and Mettel are examples of Virtual Network Operators (VNOs) that buy wholesale network access and can turn around and provide an enterprise a single entity to work with for ordering, installing, support, and billing. Large enterprises can buy wholesale bandwidth at large carrier-neutral colocation providers like Equinix, CoreSite, and Cologix, including wholesale access to large cloud providers (IaaS, SaaS, UC/CCaaS). This is a model wherein an enterprise goes to the ~1,000 fiber networks versus bringing in one or two network service providers.
2) **Be Transport Service Agnostic** – SD-WAN software provides an overlay for routing and segmentation that allows for any-to-any connectivity over a bunch of different transport types. Ethernet Private Line (EPL), VPLS, MPLS, Internet, LTE, satellite, and fixed wireless/5G are all options depending on the size and reliability requirements at a site. Most applications that run the business are lower bandwidth transactions. High-bandwidth applications are those used to manage the business — typically collaboration applications. Some retail stores run on LTE or satellite after a fiber-seeking backhoe has cut the primary link. Also, for high-bandwidth sites, EPL is typically 25% less expensive than MPLS.

3) **Commodity Hardware & Software** – Thanks to Intel DPDK and AES-NI as well as initiatives from the opensource community such as FD.io, the cost of WAN hardware and software is a fraction of what it used to be. In the branch office, the router and security software can run on the same hardware as the print server and other services. The need to have specific hardware just for WAN routing has gone away. This concept is hard for many enterprises organizationally because of the divide between who owns/manages the hardware versus the services. Giving the network team the ability to own and manage the hardware is one reason why 90% of early SD-WAN implementations are on appliances. Many of the new SD-WAN vendors have the foundation of their software with opensource components and have pricing and business models that are 75% less than the traditional router. Same goes with support costs.

4) **Automation** – Zero touch provisioning and automating network changes helps reduce the number of people required to install, manage, and support a WAN. One thing that most SD-WAN platforms do well is enable dynamic traffic routing over the best path, and when congestion or outages occur, traffic is automatically routed to a better path.

Enterprises should take a page from the cloud providers' book and follow these principles in getting faster, better, cheaper, and more secure solutions versus having to make a trade-off. The days of buying technology and services from just a few well-established vendors are coming to an end.

Cloud providers use the above principals to build their WANs. Enterprises should take a page from the cloud providers' book and follow these principles in getting faster, better, cheaper, *and* more secure solutions versus having to make a trade-off. The days of buying technology and services from just a few well-established vendors are coming to an end.

SIP trunking uptake was initially slow since there were no turnkey solutions and companies had to add a new SBC vendor, and carriers offering the biggest savings were the tier 2 and 3 carriers. The SD-WAN market is similar today: the biggest savings are not coming from the well-established vendors…yet.

Zero Trust Security to make your Network More Secure

With Zero Trust Networking and Secure Access Service Edge (SASE) becoming more important, many parties are interested in how to incorporate these concepts into their open and modular SD-WAN solution.

The basic Zero Trust Networking concept is that we stop malicious traffic before it even gets on the IP network. In this world of mobile users, billions of connected things, and public cloud applications everywhere — not to mention the growing sophistication of hackers

and malware — the Zero Trust movement is the new reality. As the name suggests, Zero Trust means no trusted perimeter — everything is untrusted and, even after authentication and authorization, a device or user only receives least privileged access. This is necessary to stop all these potential security breaches.

Zero Trust networking (ZTN) is the application of the zero trust principles to enterprise and government agency IP networks. Among other things, ZTN integrates IAM into IP routing and prohibits establishment of a single TCP/UDP session without prior authentication and authorization. Once a session is established, ZTN ensures all traffic in motion is encrypted.

To put in context of a common analogy, think of our road systems as a network and the cars and trucks on it as IP packets. Today, anyone can leave his or her house and drive to your home and come up your driveway. That driver may not have a key to get into your home, but he or she can case it and wait for an opportunity to enter. In a Zero Trust world, no one can leave their house to travel over the roads to your home without prior authentication and authorization. This is what's required in the digital, virtual world to ensure security.

In the voice world, we use signaling to establish the authentication and authorization prior to connecting the call. In the data world, this can be done with TCP/UDP sessions, and in many cases, in conjunction with Transport Layer Security, or TLS. The problem is that IP routing hasn't evolved since the mid '90s. IP routing protocols such as Border Gateway Protocol are standalone; they don't integrate with directories. Network admission control (NAC) is an earlier attempt to add IAM to networking, but it requires a client and assumes a trusted perimeter. NAC is IP address-based, not TCP/UDP session state-based.

The solution is to make IP routing more intelligent and bring it up to the OSI stack to Layer 5 where security and session state reside. The next generation of software defined networks are taking a more intelligent approach to networking with the Layer 5 security and performance functions.

Over time, organizations have added firewalls, session border controllers, WAN optimizers, and load balancers to networks for their ability to manage session state and provide intelligent performance and security controls required in today's networks. For instance, firewalls stop malicious traffic in the middle of a network and do nothing within a Layer 2 broadcast domain.

Every organization has directory services based on IAM that define who is allowed access to what. ZTN takes this further by embedding this information into the network and enabling malicious traffic to be stopped at the source.

Another great feature of ZTN is anomaly detection. When a device starts trying to communicate with other devices, services, or applications to which it doesn't have permission, an alert can be generated. Hackers use a process of discovery, identification, and targeting to break into systems; with Zero Trust, you can prevent them from starting the initial discovery. For more information on ZTN, see the "Zero Trust Networking" report published in 2018. We expect an update on this subject later this year.

Existing network approaches do not provide the levels of security and access control digital enterprises require. There is a demand for immediate access for users, no matter where they are located or which device they are on, in a way that meets all the security requirements. For too long, network routing and security have been separated. Merging them together in an open and modular way creates synergies over and above stacking the two together using Network Function Virtualization.

> When a device starts trying to communicate with other devices, services, or applications to which it doesn't have permission, an alert can be generated.

One of the advantages of SASE is modularity and using only the security components required instead of the entire security stack that comes with today's next generation firewalls. For instance, an enterprise can provide internet off-load at a campus location and create a policy that if the site is on the whitelist of approved sites and the application is TLS authenticated and encrypted with a validated certificate, then the user can route directly to the application such as Office365 or WebEx. All other traffic will then be directed to a more robust security stack that provides web filtering, sandboxing, DNS security, credential theft prevention, data loss prevention, and next-generation firewall policies.

Open and modular routing and security capabilities will enable:

1. **Flexibility**: Using what is required versus a full bloated software stack.
2. **Cost savings**: Instead of buying and managing multiple point products, utilizing a single platform will dramatically reduce your costs and IT resources.
3. **Increased security**: A Zero Trust approach to the cloud removes trust assumptions when users, devices, and applications connect. A SASE solution will provide complete session protection, regardless of whether a user is on or off the corporate network.
4. **Data protection**: Implementing data protection policies within a SASE framework helps prevent unauthorized access and abuse of sensitive data.

SD-WANs are built on overlays such as IPsec in order to get an IP packet to route across a path that the native/original IP header cannot do along with providing path security via encryption. VxLAN is another overlay used by some SD-WAN vendors to provide segmentation and encapsulation over and above what one can do with a standard IP packet.

VxLAN offers a hierarchal, end-to-end method to segment network traffic to provide the performance and security controls that digital enterprises are demanding. While there is no overall SD-WAN protocol standard, VxLAN is an industry standard that can be used in data centers, cloud providers, campus, branch-office, and VPN solutions.

Why is VxLAN important? Two of the core reasons follow:

1) **Scalability** – Traditional VLANs only scale to 4,096 unique networks within a domain. With Zero Trust requiring 1:1 micro-segmentation between users, devices, services, applications, and data, traditional VLANs do not scale. Security these days is not just about north/south segmentation, but east/west and micro-segmentation. VxLAN scales to 16 million unique networks within a domain.
2) **Blending Virtual & Physical Networks** – The VXLAN VTEP can be implemented in both virtual and physical switches allowing the virtual network to map to physical resources and network services. VXLAN Tunnel End Points (VTEP) which perform the encapsulation/de-encapsulation.

The "secret sauce" in utilizing VxLAN to provide a ZTN/SASE will be mapping Identity and Access Management (IAM) directories to VxLAN. Directory Enabled Networking (DEN) has been around for decades, but has never taken off, in part because of scalability challenges, which is the same reason routers do not manage session state like firewalls do. But as networking and routing move to all software and platforms can scale horizontally, scalability is no longer an issue.

SD-WAN Business Drivers

The digital enterprise is taking advantage of IoT, AI, Video Everywhere, and social enabling greater employee productivity, automation, efficiencies, and speed to market. Besides making WANs faster, better, cheaper, and more secure, the SD-WAN market is being driven by an explosion in applications, video, big data, and more security. Figure 12 below summarizes these drivers.

Besides making WANs faster, better, cheaper, and more secure, the SD-WAN market is being driven by an explosion in applications, video, big data, and more security.

Application Explosion	• Cloud computing with applications distributed everywhere • Mobility and consuming applications from anywhere • Application performance can directly impact business performance
Video	• Voice & video embedded into composite applications • Video using 95% of WAN bandwidth: digital signage, surveillance, video conferencing, collaboration, training, marketing, real-time streaming
Big Data	• IoT - gathering and analyzing data from millions of sources • Real-time analytics to improve business outcomes • Large data file transfers for synchronous replication
Security	• Segmentation – Guest Wifi, Credit Card, Critical Data, Voice/Video, etc. • Encryption – All data in motion • Analytics – Understanding what is going across the network at all times

Figure 12: Summary of SD-WAN Drivers

Most SD-WAN deployments are part of a greater enterprise digital strategy and overhaul of the branch offices. One very large retailer had the goals of using SD-WAN to reduce WAN costs by $100 million, reducing staffing costs by $400 million through automation and sharing resources across sites, and growing top-line revenue by $3-5 billion by offering new services such as a store within a store, video kiosks, and tele-health.

The Reality of SD-WANs

Many of the concepts underpinning SD-WAN (such as encryption, path control, overlay networks and subscription-based pricing) are not new. However, SD-WAN essentially wraps these technologies together and presents them to enterprises as a new integrated offering.

If you look at Budweiser beer can labeling over time, you can see that the label has changed but the beer has stayed the same. Similarly, the SD-WAN market started in the mid-2000s with companies like Talari coming to market with solutions to make IPsec networking better, especially for international WANs using the Internet as transport.

In 2012, the SD-WAN market started taking shape and some of the largest SD-WAN vendors started in this timeframe such as Viptela and Velocloud. In 2015, SD-WANs were at the peak of

the hype cycle or as others would say, "inflated expectations." Today, we are coming out of the trough of disillusionment as enterprises learn the hard way that SD-WANs are complex and that there is no one-size-fits-all solution.

The evolution of the WAN can be seen in Figure 13 below. While many will say that SD-WAN is a revolution, the reality is that the technology itself is part of the WAN evolution. The revolutionary part is how enterprises can use the technology to help transform their businesses.

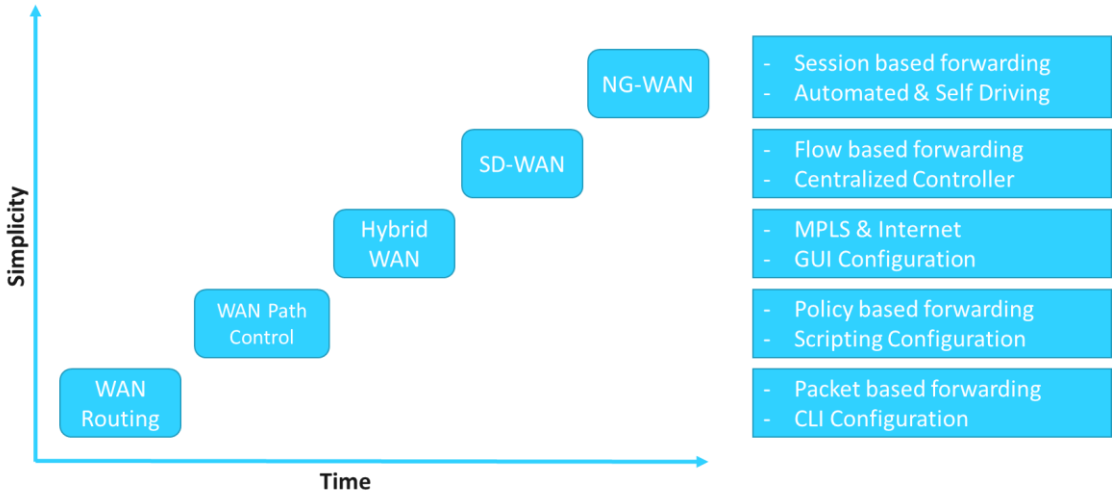

Figure 13: Summary of WAN Evolution

IPsec tunnels were the start of creating SD-WANs by getting network to route packets via a path that the original IP header would not route over and to provide encryption for data in motion. Because IPsec has limitations, new vendors came along and started adding features such as bonding multiple IPsec tunnels together, providing brown-out routing, adding segmentation, and many other features over and above what IPsec natively offers.

SD-WANs are getting a lot of good press. Here are issues to be aware of when building an SD-WAN strategy and selecting an SD-WAN vendor.

No Standard – and This is a Big Problem

There is no SD-WAN protocol standard, and this is a big deal! Unlike all other network technologies that have defined protocols by the IETF, SD-WAN does not have a defined protocol standard. No SD-WAN standards means enterprises and service providers will be locked-in to the SD-WAN vendor(s) that they choose. Also, no cross-vendor interoperability, no ability to troubleshoot problems against a known standard. If a vendor goes out of business or gets

No SD-WAN standards means enterprises and service providers will be locked-in to the SD-WAN vendor(s) that they choose.

acquired by another vendor and the talent leaves, enterprises and service providers are at risk of having to deploy another solution down the road.

The promise of Software Defined Networks is to enable networking to be hardware, transport, service provider, and vendor agnostic. Part of this open architecture was separating the data, control, and management planes and becoming a pure software solution which could be deployed anywhere and everywhere. As the networking industry moves towards this objective, we are lacking in one key area, which is the creation of a common SD-WAN overlay protocol.

The challenges of today's SD-WAN environment is an enterprise or service provider must select a specific vendor or two, and then suffer the risks and costs associated with SD-WAN vendor lock-in. Every SD-WAN vendor uses a proprietary label as a header to add to every native IP packet. This label is composed of IPsec plus other headers including VxLAN for segmentation, and other proprietary headers to improve network security and performance. Below is an example of an SD-WAN packet with IPsec & GRE overlay headers.

Figure 14: Example of the Overhead Added to Each IP Packet in SD-WAN

One example of how this overhead is really inefficient is with a G.729 call that sends a 20byte data packet of the voice sample every 20 milli-seconds. Without SD-WAN, it would be a 60byte packet with the original IP header of 20 bytes, and 20 bytes in UDP/RTP. The GRE overhead is 24 bytes and the IPsec overhead is 50 bytes. So, a 60-byte voice packet grows to 134 bytes.

None of today's SD-WAN headers are compatible, forcing enterprises and service providers to backhaul traffic to a communications hub to internetwork with their non SD-WAN networks, and in the process lose all the enhanced security and performance features that SD-WANs provide above standard IP/BGP networking.

To date, many have tried, but all have failed to create an SD-WAN overlay protocol standard. A few noteworthy attempts to standardize SD-WAN include:

a. **ONUG** – Gave up at the protocol level, trying at the orchestration - OSE
b. **Open Networking Linux** – Goal is open source networking, not SD-WAN - OCP
c. **IETF** – Yang VPN standardized SD-WAN descriptions and overall SDN standards.
d. **MEF** – Certification, coordination of SDN, NFV, VNF and open source networking. MEF 3.0 SD-WAN standard describes requirements for an application-aware, over-the-top WAN connectivity service that uses policies to determine how application flows are directed over multiple underlay networks irrespective of the underlay technologies or service providers who deliver them.

The industry needs an SD-WAN protocol standard. If this cannot be done by the standards body, the other option is for a solution to become the defacto standard. Opensource options such as the one from flexiWAN is one path to this. An SD-WAN overlay protocol would be a combination of IPsec plus VxLAN, plus a new 2-byte field. The rationale for this is:

1. **Encryption & Security** – IPsec is a common standard that has been around for decades, though cross vendor implementations can be problematic. IPsec is the industry default for creating tunnels and providing encryption for data in motion over networks. This is why every SD-WAN uses IPsec tunnels as part of their overlay protocol foundation. (The reason the SD-WAN market got started back in the mid-2000s was to address the limitations of IPsec, especially in international WAN implementations.)
2. **Segmentation** – There are multiple options listed below to segment network traffic, but there is a case for VxLAN:
    - **VLAN** – Traditional layer 2 segmentation which is quite common at campus and branch sites but lacks scalability in micro-segmentation in data centers and cloud environments since there is only space for 4096 VLANs. As the world moves to a Zero Trust Networking (ZTN) model with 1:1 segmentation, VLANs have no chance of expanding into this model.
    - **GRE** – Traditional layer 3 tunneling and lacks scalability and multi-vendor supportability. Most intra-site segmentation is done at layer 2 to keep costs down and performance up. Cisco's IWAN and a lot of proprietary cloud connections used GRE and found it complex to manage.
    - **Proprietary** – More bandwidth efficient, but very little industry adoption. For instance, it appears that the Viptela uses a 4-byte proprietary tag in the SD-WAN header, of which one byte (8bits) is set aside for segmentation. This means the total number of segments is limited to 256 (2-8th)
    - **The case for VxLAN** – A modern layer 2 segmentation solution used commonly in data centers and clouds.
        - Scalability – Number of segments supported is 16 million, and these segments can be designed in hierarchies for easier manageability. RFC 7348 VXLAN each segment is identified through a 24-bit segment ID, termed the "VXLAN Network Identifier (VNI)." This allows up to 16 M VXLAN segments to coexist within the same administrative domain.
        - End-to-end – Segmentation from LAN across a WAN to the far end LAN and not having to map to layer 3 VRFs.

- Standardized – Common and widely deployed in next generation private and public data centers.
- VxLAN is used today in some leading SD-WAN solutions including VMware/Velocloud and Versa and also flexiWAN, the opensource SD-WAN solution.

- **Challenge of VxLAN** – Adds an additional 50 bytes to every packet compared to 24 bytes for GRE or 4 bytes to some proprietary schemes.

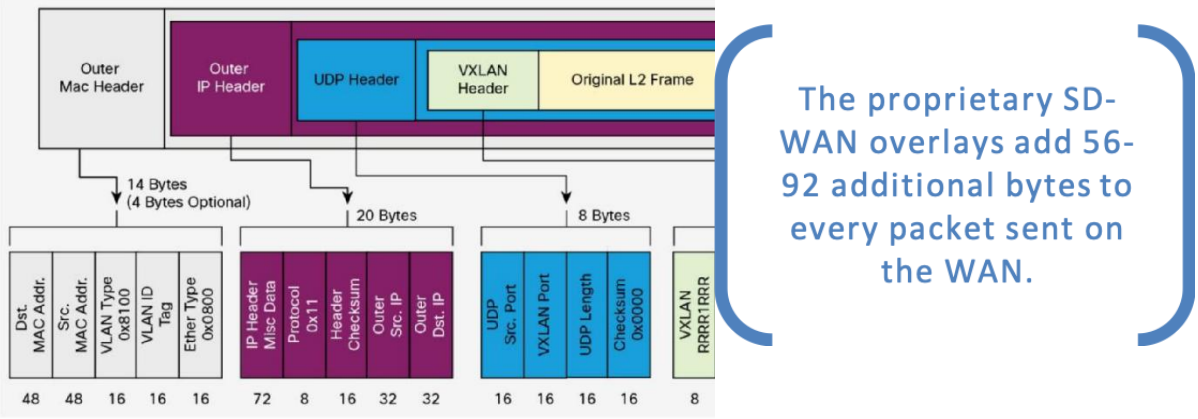

Figure 15: VxLAN Protocol Format

3. **Performance Monitoring** – This can be done by periodically sending pings or BFD to keep alive packets, putting time stamps on packets by adding this to the packet header, or creating a unique sequence number for packets and using Network Time Protocol (NTP) to track this. All of these methods work towards the same goal of allowing application traffic to run over the best path and route around "brown outs."

Creating an open and common SD-WAN overlay protocol is a win for everyone. Enterprises and service providers avoid vendor network lock-in and do not have to settle for a single SD-WAN solution for all of their sites. Instead they can choose different platforms based on different use cases that vary by costs, site size, security requirements, and specific application performance. For SD-WAN vendors, a common protocol leads to faster industry adoption with the differentiating features added higher up the stack while reusing common components at the lower part of the stack. This is part of the promise of a software digital world where everything is interconnected.

Until an SD-WAN standard exists, enterprises and service providers will have to utilize IP/BGP networks as the least common denominator to internetwork different SD-WAN solutions. This is done today with email via SMTP and other higher-level applications and is doable for SD-WAN, but far from ideal.

25 - 50% "Bandwidth Tax!"

The proprietary SD-WAN overlays add 56-92 additional bytes to every packet sent on the WAN. Today's MPLS adds 4 bytes of overhead as a tag to separate various IP networks that run over it. If the average packet size is 160 bytes, this overhead is significant. This overhead can also add to large packet fragmentation which will break many applications. This is shown in figure 15 above. Another graphical representation of this is below in figure 16.

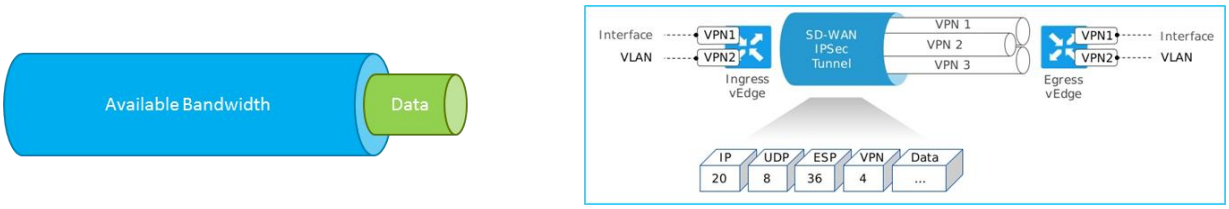

Figure 16: SD-WAN Overhead

The overhead is even greater when SD-WAN vendors add additional overhead from additional tunnels such as GRE or VxLAN. 128 Technology is the only vendor that does not use tunnels and does not add overhead to every packet unless encryption is required, since they are session aware and maintain session state in their router.

Forward Error Correction (FEC) adds additional overhead along with link and application monitoring. Some SD-WAN vendors use time stamps, or the equivalent of IP sequence numbers tied to Network Time Protocol (NTP) to monitor latency, jitter, and packet loss on a link and per application. Other vendors use Bidirectional Forward Detection (BFD) or pings to measure each link. In a mesh architecture, this can quickly add up to significant overhead.

SD-WAN vendors will show how file transfers are faster when there is packet loss on a WAN link and that FEC is worth the overhead. The problem with their simulations is that they have the packet losses occur randomly. The reality is that packet loss comes in bursts, and FEC across a single link does not help in this case. The challenge of running FEC across multiple links is adding additional buffer delay, mentioned below.

Fragmentation is another problem that adding the tunnels and labels add. For instance, a large file transfer or video conferencing session will use the maximum sized standard Ethernet frame size of 1500 bytes and then the 76-90 bytes SD-WAN overhead is added, which when transported over WAN Ethernet causes fragmentation with a second packet being created. This second smaller packet may arrive sooner than the first packet based and has to be reassembled and buffered by the far end router which adds latency and takes CPU cycles. The Path MTU Discovery protocol can be used to check the end-to-end packet size so the MTU on the transmitting device can be set appropriately and avoid fragmentation.

The quantity of management data going from the site router to a centralized management platform can consume a lot of data. In fact, in one SD-WAN deployment, the SD-WAN management data was the third highest talker on the network. Similar to how 40% of the phone bill, is the phone bill, in the SD-WAN world 25 - 50% of the bandwidth can be the SD-WAN solution itself. In the early 2000's when bandwidth was 100x more expensive than today, moving to VoIP was an issue because of the overhead that IP/UDP/RTP added.

The SD-WAN market is focused on the edge of the network with speeds up to 1Gbps per site.

Yes, bandwidth is getting cheaper, but there are many situations where bandwidth is at a premium such as using LTE, satellite, or private international links. One retailer did a file transfer test of a 6M (4x1.5 T1) MPLS link versus a 10M (Ethernet) SD-WAN Internet link. The 6M MPLS link was 20% faster, even though the bandwidth was 40% less. This is real world validation that the SD-WAN overhead is significant and impactful.

Lack of Scalability

SD-WANs use tunnels, whether it is IPsec-based or a proprietary schema. Tunnels take router resources for managing tunnel state, plus the encryption. To provide a full mesh of tunnels to all sites on a large WAN requires N*(N-1) connections. For site-to-site connectivity, most SD-WAN players will backhaul traffic to a data center and avoid the mesh. Also, every tenant, which is a local LAN segmentation, requires its own tunnel such as guest Wi-Fi, voice, and credit card traffic. A 100-site WAN with 3 WAN links per site (MPLS, Internet, LTE) and 8 tenants per site will require close to half million tunnels in a full mesh.

Similar to how 40% of the phone bill, is the phone bill, in the SD-WAN world 25 - 50% of the bandwidth can be the SD-WAN solution itself.

One of the reasons that we went away from frame-relay to MPLS/IP was to migrate from hub and spoke architectures to an any-to-any model. Hub and spoke architectures in SD-WANs is a step backwards and will need to change as latency requirements on WANs move down to 5ms between users and their applications.

The SD-WAN market is focused on the edge of the network with speeds up to 1Gbps per site. Core network backbones that run at speeds from 10Gbps-1Tbps are not a good fit for SD-WANs since the CPU and bandwidth overhead from encryption is too high. SD-WAN encryption should be done at the edge of the network, not the core. Zscaler is an example of this, where at the time of writing this research, they only support 300Mbps via IPsec tunnel or 1Gbps via GRE tunnel per site.

New Hardware Required

IPsec and other tunneling encryption protocols require significantly more CPU than existing routers can provide. SD-WAN migrations require a new hardware platform, whether it is a fixed appliance or a virtual server at the branch office. Cisco ERS, which is their software-based router that runs on Linux takes an approximately 10x hit in throughput performance when 256bit encryption is turned on.

Part of the reason that SD-WANs are a revolution is that new hardware is required. This is not a feature that you turn on in an existing router. Cisco has positioned their Viptela software to run on ISRs, but many enterprises have found that they have to buy new ISRs in order to get the throughput and stability they desire.

There are many hardware options, but using commodity off-the-shelf hardware allows enterprises to get a $750 box that can provide 1Gbps of encrypted throughput. This commodity hardware is about one-tenth the cost of what traditional dedicated hardware appliances costs. Thanks to Intel's DPDK initiative, the network throughput performance using their chipsets has vastly improved, and thus the proprietary chipsets that went into networking appliances are going away.

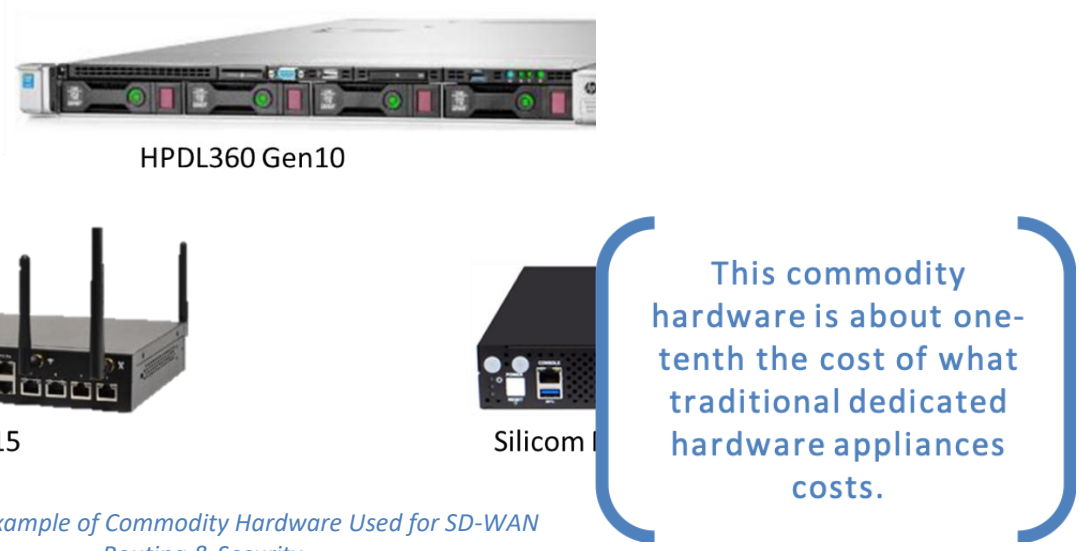

Figure 17: Example of Commodity Hardware Used for SD-WAN Routing & Security

So yes, a forklift upgrade is required for SD-WAN. The good news is that hardware is getting cheaper and the CPU cycles required to do 256bit encryption have been optimized.

Poor ROI for Large Sites

Broadband Internet access is the primary driver of savings, but broadband speeds top out around 50/5Mbps. Large sites that require 100Mbps and greater speeds bi-directionally still need to buy dedicated Ethernet access, regardless if using MPLS, VPLS, or Internet. The premium of MPLS

over VPLS and Internet is only 10 - 20% in the U.S. and even less in the rest of the world. As mentioned earlier, network transport access represents 60% of the cost of most WANs.

As broadband Internet scales up, so do the needs of sites. Most enterprises find that their WAN bandwidth needs grow about 30% per year. So as broadband Internet scales up to 1Gbps, many large sites need 10 - 100Gbps for their manufacturing, large hospital, or campus sites. For this reason, the largest adoption of SD-WAN to date are organizations that have many small branch locations with under 200 people such as banks, retailors, or clinics.

SD-WAN technology is focused today on connecting the small and medium branch office to the data center and cloud. Large campus, R&D, manufacturing, and data center and cloud interconnects, and other large sites that are running speeds of 1 Gbps or higher typically do not find benefits from SD-WANs.

60 - 80ms of Additional Latency

Many SD-WAN vendors have Forward Error Correction (FEC) via packet duplication, which they run across multiple network links. In order to provide FEC and to adjust for packet fragmentation — and the variable latency across different paths — they add a buffer to synchronize the packets, which adds delay.

It is the author's opinion that FEC and encryption should be done end-to-end, not in the middle of the network with SD-WAN which is a point-to-point connection. Modern codecs have FEC built into them and TLS encryption. FEC is something that makes sense in specific situations, but in most cases, it does more harm than good in terms of increasing latency and decreasing bandwidth.

Inefficient Double Encryption

Many applications such as virtual desktops are encrypted at the application layer using TLS, and do not need to be encrypted a second time at the network layer, which adds additional network overhead and router CPU utilization. Of the traffic on the Internet, 75% is TLS encrypted. In the near future, all applications on both public and private networks will be TLS encrypted.

Doing IPsec encryption on top of TLS encryption adds little value. One of the big differences between TLS and IPsec encryption is that TLS encrypts the data after the TCP header whereas IPsec encrypts the TCP header. With TLS, WAN optimization can manage things like the TCP window size and TCP retransmits

Some of the leading SD-WAN vendors are starting to provide adaptive encryption where they can detect if an application is TLS encrypted or if a session needs to be encrypted versus just encrypting everything regardless if the session needs it or not. This is a more dynamic and intelligent way of providing encryption.

Too Many SD-WAN Vendors

There are over 60 SD-WAN vendors. Many startup vendors have been acquired by much bigger companies, and many non-WAN companies are pivoting into SD-WAN to try and take advantage of the 70% CAGR in this market. These vendors fall into six categories:

1. **Traditional** – Vendors that have been in WAN routing for decades include Cisco, Juniper, and Huawei.
2. **Pureplay** – Vendors that are just SD-WAN solutions and where startups focused in the market include: Talari (owned by Oracle), Viptela (owned by Cisco), Velocloud (owned by VMware), Versa, and Cloudgenix. Many of these startups came from those who worked at traditional vendors, who could see that the vendors weren't pivoting fast enough for the next generation of WAN technologies.
3. **WAN Op Pivot** – Vendors that were in WAN Optimization or other networking function include Riverbed, Silverpeak, and Citrix. These vendors have an installed base that they can upsell SD-WAN into.
4. **Security Pivot** – Vendors that make firewalls and other network security products include Barracuda, Fortinet, and PaloAlto. These vendors have an installed base that they can upsell SD-WAN into.
5. **NaaS** – Network as a Service vendors offer both edge hardware and software along with providing a global network to help avoid some of the challenges within the 2,000+ Internet exchange points worldwide. NaaS vendors include Arayaka and Cato.
6. **Disruptors** – Vendors that are changing the way IP networking works and creating new routing protocols and security models include 128 Technology and Tempered Networks.

Since no one vendor has over 25% of market share and there are so many vendors, the process of deciding which vendor(s) are the best for an enterprise is complex.

There is a lot of risk in betting on a vendor, since small vendors that are easy to work with can get acquired by big vendors with complex processes, or they can simply go out of business. The price difference for hardware, software, and support between a traditional vendor and early-stage vendor can be a 10x difference. There are enterprises that have put in new SD-WAN hardware and software that is 10x faster and are paying less per month than they were under their old support contracts for older hardware and software.

Below is a sample list of vendors and logos in no particular order.

Figure 18: Examples of SD-WAN Vendors

Moving to SD-WAN is still compelling for many enterprises, even with all of these limitations. The problem with SD-WANs is they have been over-hyped, and the goal here was to outline the reality of what they do. Since every WAN is unique, the next section covers the major architecture considerations that go into a design that leads to a great solution.

Architectural Considerations for SD-WAN

Most SD-WAN deployments are part of a greater Digital Enterprise strategy that often includes an overhaul of the branch offices. New requirements driven by new applications and business objectives require the WAN to be architected differently than it is today. As mentioned earlier, the number of architectural considerations vary widely and leads to SD-WAN being implemented with very different designs and operational models. We'll now examine these architectural considerations starting with a top down model.

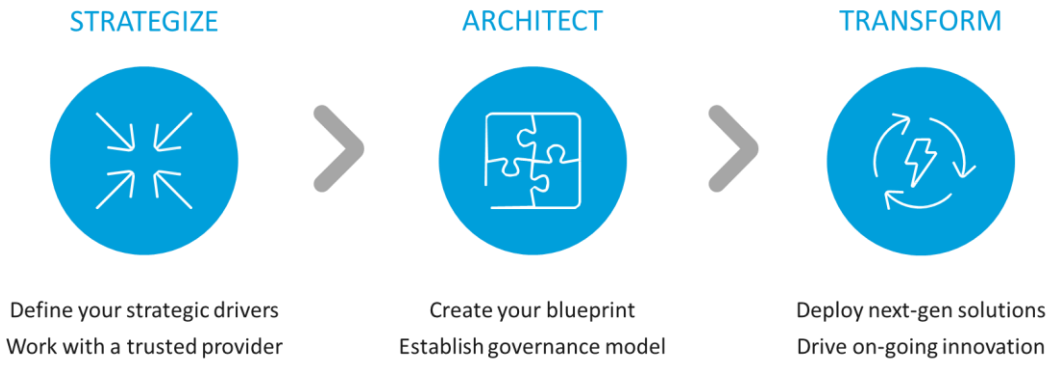

Figure 19: Strategy to Transformation Model

Figure 19 above shows the standard model we've worked with clients on in major projects that includes defining strategic business drivers, architecting a solution, deploying the solution, and transforming the business. Because there are a significant number of architectural considerations, we'll describe the various architectural areas and how they fit.

Business Architecture

Business architecture represents strategic objectives in creating, selling, and supporting products and/or services including business processes, information, and governance. It is used to create competitive advantages, leverage existing strengths, and identify investment opportunities that advance the business's strategic objectives and drive innovation. This is then used to develop plans, make business decisions, and guide implementations.

Architecture is about making decisions when there are competing requirements of finding the right balance of faster, better, cheaper, and more secure. Below are some common business architecture decisions that need to be made:

Business Metrics and Defining Success

Happiness is reality minus expectations. Setting expectations and measuring the reality to those expectations will define a successful project. Many business leaders expect their corporate networks to be another service like cloud services that are on-demand and elastic, where you pay for what you use. The concept of pre-paying or getting into long-term contracts is no longer in vogue.

New expectations are for push button infrastructure, including networking. It should be this simple, easy, reliable, and cheap with great security built in. The engineering reality of how a solution works can be vastly different from the product and marketing. Common metrics that need to be defined fit into the following buckets:

1) **Faster**
    a. **Network Performance** – Salesforce application login screen will come up in ¼ second. A 10M file will download from MS Office 365 email in 2 seconds.
    b. **Agile** – Add a new site in five business days, and ability to scale up or down bandwidth on demand and within five seconds.
2) **Better**
    a. **Automated** – Ability to turn up a new site without dispatching a technician and add new network features or upgrades with a click of a button.
    b. **Integrated** – DevOps tools so that networking bandwidth and associated routing and security policies can be provisioned dynamically as compute and storage change. Add a new VM and associated network in two seconds.
    c. **Highly Reliable** – 99.999% reliability where the network downtime to a site is less than six minutes.
3) **Cheaper**
    a. **Lower operating expense** – Cut networking costs by 30% - circuits, staff, maintenance
    b. **Lower capital expense** – Cut hardware and software costs by 75%
4) **More Secure**
    a. **Encrypt All Data in Motion** – Everything is encrypted everywhere, even on the network.
    b. **Least Privileged Access** – Create a whitelist of what users, devices, services, applications, and data are allowed to talk with instead of a blacklist of what they are not allowed to talk with.
    c. **Minimize the Attack Surface** – Score of one for everything (see ZTN research for details).
    d. **Identify Issues in < 2 seconds** – Anomaly detection of malicious users or malware in seconds.

The old saying of "if you cannot measure it, you cannot manage it" applies here. Furthermore, if expectations are defined up front, continued measurements and reporting will go a long way to success. Going from 99.95 to 99.999% reliability can have a significant impact on costs, and these trade-offs are part of the business architecture process.

Who Controls the SD-WAN solution – IT or Business?

Shadow IT has emerged as business leaders take direct control of their applications, and thanks to the cloud, they are less dependent on IT for the infrastructure. The network is the last thing

that IT fully controls within the enterprise, and it consumes 10% of the enterprise technology budget. More and more technology decisions are being made by business leaders who focus on results and don't want or need to be involved in the underlying solutions and vendors.

Users can get frustrated with their enterprise network because it is often slower to work in the office than it is to work from home, and this seems counterintuitive to both the employee and executives. CIOs wonder why they pay 20x more for enterprise bandwidth than what they pay as a consumer. Business leaders are also frustrated with the enterprise network because it is slowing down their digital transformation projects.

Enterprise networks are inherently slower, less agile, less secure, and more expensive because of:

1) **Backhauling** – Sending all Internet destined traffic back to a data center before going out to the Internet. Eight percent of enterprise branch office traffic is Internet destined and the backhauling is both expensive and slows down cloud-based applications. Mobile device managers also backhaul cellular data traffic, causing the same problem.
2) **Legacy business models** – Buying upfront tons of equipment (routers, firewalls, load balancers, network optimizers, intrusion detection) and signing multi-year contracts with one to two network service providers.
3) **ACLs are out of control** – Access Control Lists are used by network equipment to define on every interface where packets can and cannot go. This manual process can lead to thousands of rules, and things spiral out of control with no one understanding why a rule put in three years ago still applies. Also, routers are not able to report on which ACLs are used. Every network change requires new ACLs, which can break existing applications, making networks very complex and fragile.
4) **Perimeter Security** – The assumption that a private network is more secure has not proven true as many hacks have been published at a greater frequency. A Zero Trust model is required to provide end-to-end security.

SD-WANs are just a step towards the Next Generation WAN (NG-WAN) that will be managed by cloud providers through Network as a Service (NaaS). Microsoft, Google, and other large Cloud Service Providers (CSPs) are becoming network operators. Gartner reports that 50% of cloud implementations have business impacting problems due to the network. CSPs realize that if they are going to provide a Quality of Experience (QoE) for their applications, they need to have greater control of connecting their users.

To achieve complete end-to-end control of business IT computing and migration to cloud services, CSPs will offer secure seamless networking solutions to connect from customers' on-premises servers to in-cloud-based resources. The next generation networks will leverage broadband Internet connectivity, high speed optical, and Ethernet networks that are inter-connected at the carrier neutral collocations where the CSPs reside. On the premises will be white box switches and wireless local area networks connected to a very intelligent router and

security stack that can dynamically establish direct, secure sessions between application services and users.

This can be done at a fraction of the cost because CSPs already possess significant technical resources in networking, and they have different business models than the traditional Network Service Providers (NSPs). CSPs over time will marginalize existing NSPs and shed the complexity that inhibits broader migration to cloud-based services.

The market for enterprise networking will go through a radical shakeout and will become commoditized. Whitebox providers that develop the appropriate partnerships will see new opportunities. Winners will include low-cost access and transport service providers along with existing and new network equipment providers bold enough to morph into a volume player for a low-margin business.

The best lens into the IT future is to watch what start-up companies are doing. These companies do not have any legacy baggage, and they adopt the latest and greatest technology and solutions. Few start-ups are creating their own private networks. AirBnB and Uber are examples of companies without a private MPLS WAN. These companies use IaaS & SaaS, and they are stitching it together with cloud Networking as a Service (NaaS).

This is a paradigm shift for the enterprise to go to the 1,000-plus fiber networks and Internet Service Providers (ISPs) that cloud providers use, versus bringing one to two NSPs and ISPs into the enterprise. Since business leaders most often pay for the cloud services that they use, the network will become an extension of this model.

To answer the question, who controls the SD-WAN solution — IT or Business? The answer is the business will own and control their network, and IT will control the interoperability between various business networks. Just like the Internet is a network of networks, enterprises will be the same way, and a smaller IT team will focus on the Internetworking.

What Enterprises can do to Save Money & Improve Performance

Enterprise network bandwidth requirements double every 2.6 years and the number of devices on the network doubles every 1.8 years, which keeps network managers and engineers busy just keeping up. Most enterprises look to get the best discounts for what they buy, but they end up saving pennies on the dollar compared to starting from scratch with a new architecture and sourcing model. Smart enterprises can cut spending by up to 50%, starting with these three strategies:

1) **Buy Wholesale** – Instead of bringing one or two Network Service Providers (NSPs) into the enterprise and paying retail rates, enterprises should go directly to the 1,000-plus fiber providers. This can be accomplished two ways. First is providing connectivity into a

carrier neutral co-location provider, such as Equinix. The alternative is to go with a virtual network operator that buys wholesale network services and provides a managed service such as MetTel.

2) **New Business Models** – Instead of paying for hardware/software upfront, pay for consumption and value. One reason the cloud is taking off is the operating expense and success-based business model. All networking and communications are going to software that runs on commodity off-the-shelf hardware. There is no need to have to pre-pay for this software. Enterprises should also negotiate with their vendors to have maintenance start when the service goes live, not the day it is ordered. Paying for maintenance on legacy gear is also a waste of money. E-Bay and other websites sell legacy hardware and organizations can self-spare. Enterprises self-insure on healthcare, and they can do this on IT hardware.

3) **Right Architecture** – Too many enterprises backhaul their Internet destined traffic through their data centers. This private-to-public Internetworking is expensive and hurts performance. An enterprise that has the best MPLS and Internet retail rates but has an architecture where all branch traffic must route through a data center to go to the Internet, are paying for bandwidth twice. Most enterprises see upwards of 80% of their branch traffic is Internet destined. The same holds true for mobile data going through a mobile device manager that is in the enterprise data center. This is why SD-WAN is forecasted to have a CAGR of 70% because it empowers local Internet offload.

To improve network performance, reliability, agility, and security while cutting network costs requires a totally different approach than what traditional carriers and vendors sell. Cloud providers have adopted the above strategies along with leading enterprises who have a proven track record that getting more out of your network for less is possible. While it is easy to spend money, it is hard to save it!

Adopting a Multi-Vendor SD-WAN Strategy

2020 is the year that SD-WAN implementations are set to explode. That said, many enterprises are nervous about their SD-WAN vendor selection, as this is a new market with a lot of new vendor and products. For many enterprises, a multi-vendor SD-WAN strategy may be the best near-term strategy as this market continues to mature. This will resonate well with enterprises that have a "best of breed" versus a "one vendor for all" strategy. All products and technologies have their time and place. Trying to force a "one for everything" model is where IT has traditionally gotten itself in trouble and one of the reasons business leaders despise IT.

Every SD-WAN solution is proprietary as was covered earlier. Cross-platform interoperability will not occur in the foreseeable future. Cisco's Meraki, Viptela, and IWAN products, for example, will never be able to directly peer with each other. Fundamentally, each uses a different "label and tunnel" format on every IP packet that is sent. That said, one of the benefits of SD-WAN platforms is simpler management, so supporting a multi-vendor or multi-product strategy is not overly burdensome.

A multi-vendor approach goes against everything that IT is about in trying to simplify the infrastructure. But this is why it makes sense:

1) **Economics** – Large enterprises divide their WAN into five different site models based on the number of users and devices at the site, as well as its business importance. Economically speaking, putting a pair of expensive routing hardware and software licenses at the smaller-tier sites does not make financial sense.

2) **Speed of Implementation** – The economics of SD-WAN for small sites is a slam dunk. Delaying the rollout while trying to get a single enterprise SD-WAN solution does not make financial sense. Plus, some SD-WAN "lite" solutions — for small sites with few requirements over and above the basics — allow for quick rollouts. Zero Touch Provisioning (ZTP) is the hot buzzword for enabling a quick, low-touch implementation model.

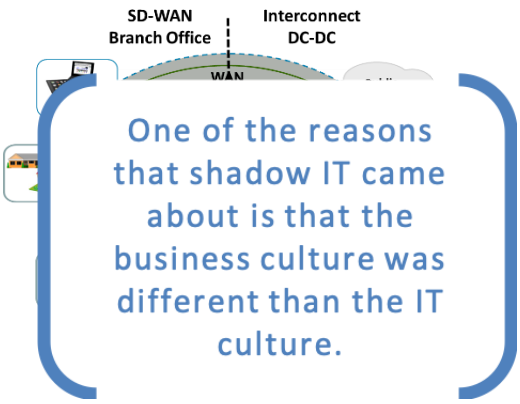

3) **Different Requirements** – The WAN requirements for connecting sites with users to data centers and clouds where applications reside is different than inter-connecting data centers and cloud services. Bandwidth, latency, segmentation, WAN optimization, and load balancing requirements are very different in these networks.

The greatest disadvantage of moving to a multi-vendor SD-WAN strategy is continuing with a hub-and-spoke WAN architecture, where all network traffic transits through a data center. While most large enterprise and government networks are designed this way today, the long-term strategy is to move to a peer-to-peer edge network to reduce network latency. Augmented reality and edge computing to support applications such as manufacturing reliability and driverless cars require users and applications to be within 5ms transport of each other.

Most enterprises support many different types of operating systems, databases, and security solutions. Not having a one-size-and-vendor-fit-all solution for network routing in the enterprise WAN will become the new norm, especially as networking moves to all software that can run on commodity and virtualized hardware. In an ideal software world, if you do not like one SD-WAN vendor or product, just delete the software stack and add the new one, re-using the hardware.

Business and IT Culture and Guiding Principles

The business and associated technology culture of large enterprises is vastly different, even if they are in the same business vertical market such as retail. One of the reasons that shadow IT came about is that the business culture was different than the IT culture. These culture

differences and associated guiding principles, that may or may not be documented, should be understood before embarking on any large IT project including the migration to SD-WAN. A sampling of these differences includes:

- **Buy vs. Build** – Data centers, networks, applications such as CRM and ERP are things enterprises can build themselves or buy/lease from others. If an enterprise decides to move to a managed network service for their SD-WAN implementation and they are used to building solutions themselves, everyone should be aware that this is a big culture shift.
- **Opensource vs. Vendor Provided** – Opensource in the networking field is just starting to take off, but is very popular in other parts of IT, with many enterprises having an "opensource first strategy." The bulk of SD-WAN solutions are built on the foundation of opensource which in part has enabled so many vendors to get into the market. When enterprises that have relied on Cisco in the past go to a new vendor or to flexiWAN, the first opensource SD-WAN solution, everyone should be aware that this is a big culture shift. Engineers have been known to leave their jobs based on vendor decisions.
- **Brute Force vs. Tightly Manage** – Overbuilding the WAN versus using Quality of Service (QoS) to tightly manage the limited amount of bandwidth used. Enterprises that typically have many small branch offices tend to tightly manage their bandwidth whereas enterprises with large manufacturing and R&D centers tend to overbuild their network. An overbuilt network does not need all the features that many SD-WAN vendors offer. The large cloud providers tend to overbuild their networks. Most of the earlier SD-WAN adopters were companies that tightly manage their WAN, so deploying an SD-WAN solution into an overly built WAN may not have as many benefits.
- **Early vs. Late Adopter** – Enterprises that adopt technology early see business value in doing so and staying ahead of their competition. Late adopters tend to see technology as cost of doing business and will provide just enough money and resources to meet the base business needs. Early adopters of SD-WAN are not only looking at lowering bottom technology costs but rolling out new applications to improve productivity and grow top line revenue.
- **Minimal vs. Maximum Security** – Many enterprises operate in regulated markets that they must comply with. Security, compliance, and privacy are top of mind these days but that does not always translate into ensuring that security requirements are the top priority. Zero Trust Network security is a new security paradigm that not all enterprises are ready to embrace.
- **Good Enough vs. Great User Experience** – Large projects many times do not specify requirements over and above getting it to work. This good enough attitude and culture in IT can frustrate business users who have a better user experience outside of the enterprise's network. Network designs that do not take into account the latency between users and the applications they consume represent a "good enough" culture. Part of the value proposition of SD-WAN is better application controls on the network and delivering a great user experience.

- **Centralized vs. Distributed Managed** – Who owns and pays for the network: Is the network run as a shared utility centrally controlled within IT, or does each business unit own and control their network? The centralized versus distributed debate has been going on for decades, swinging back and forth between the two. A single versus multi-vendor approach is a subset of this. End-to-end versus a point solution applies to an SD-WAN design that may just focus on connecting the branch office to a data center and not users to applications.

How one approaches the selling, design, and implementation of SD-WAN should take these cultures and IT guiding principles into account. What works at one large retailer may not work at another because of the different business and IT cultures.

Managed Service Versus Do It Yourself

Most enterprises are choosing a managed service provider to implement and support their SD-WAN. As more business leaders make technology decisions, they are also deciding to go with a managed SD-WAN solution and be less dependent on IT.

Traditionally, 35% of U.S. enterprises used a third-party managed service for their WAN. In the rest of the world, it is around 80%. In the SD-WAN market, 80% across the board are going with a managed service offering, and this number is expected to go up. The top reasons for using a managed service are:

1. **Implementation Speed** – As discussed earlier, most WAN refreshes take years. The ROI and business benefits that SD-WANs offer is so compelling, organizations want to get them done in 6 - 9 months. A turn-key solution with cloud-based management really reduces implementation time.
2. **Experience** – SD-WANs are new, and the skills and experience are lacking in the market. Since every solution is proprietary, one needs the skillsets from the vendor(s) chosen.
3. **Management** – Day two operational support and the ability to make changes quickly while mitigating risks. A great portal shows the health and utilization of the network and applications riding on it, including voice quality metrics.
4. **Single "Throat to Choke"** – Since SD-WANs can be used by many different underlying network service providers, having single ordering, provisioning, support, and billing is needed.

Enterprises that choose a managed network service do so for the reasons stated in figure 20.

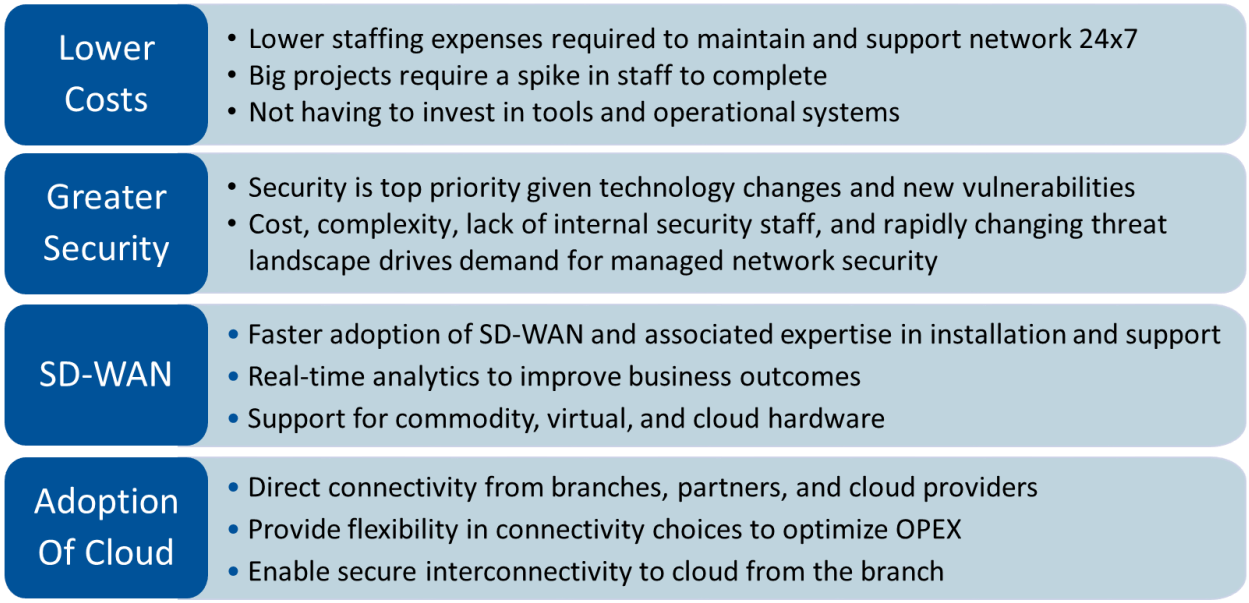

Figure 20: Advantages of Managed Network Service

Once the business architecture is done, it helps feed the technical architecture. Just like designing a new home, one must start with what the clients wants and can afford. There are many regulatory requirements just like there are building code requirements which all factor into the solution. Too often a vendor comes in and sells a solution, and the technology drives the architecture versus the business. Many business leaders know what they do not want, before they see and realize something that they like. Having a managed service provider share some of the business and technical architecture and implementations that they have already done elsewhere goes a long way towards setting expectations.

Technical Architecture

Technology architecture is the development of design and guidelines within an architectural framework. Using best practices and technology standards, a blueprint is built showing current and future state and the major steps required for the transition. A lot of variables go into building a technical architecture.

As covered earlier, WANs are built on access, transport, and hardware/software. SD-WANs add an overlay on top of this, then we cover security and unified communications separately because they are core to the overall success of an SD-WAN architecture.

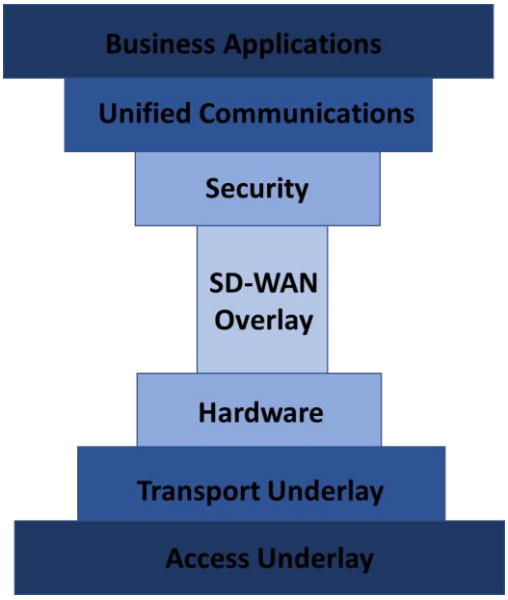

Figure 21: Building Blocks of an SD-WAN

Access Circuits – The Network Underlay Foundation

The fact that 60% of WAN costs and 95% of problems reside in the last mile of network access, means that the access strategy is core to any WAN architecture.

Network access is the connection from a building to a network service providers (NSP) or a network collocation facility. Access is focused on layer 1 of the OSI model on connecting one location to another.

A good access strategy is critical, with more enterprises leveraging a combination of wireline and wireless access to improve reliability. Public and private 5G will further enable businesses to use wireless as a high-speed WAN access strategy. Network slang for access is the "last mile."

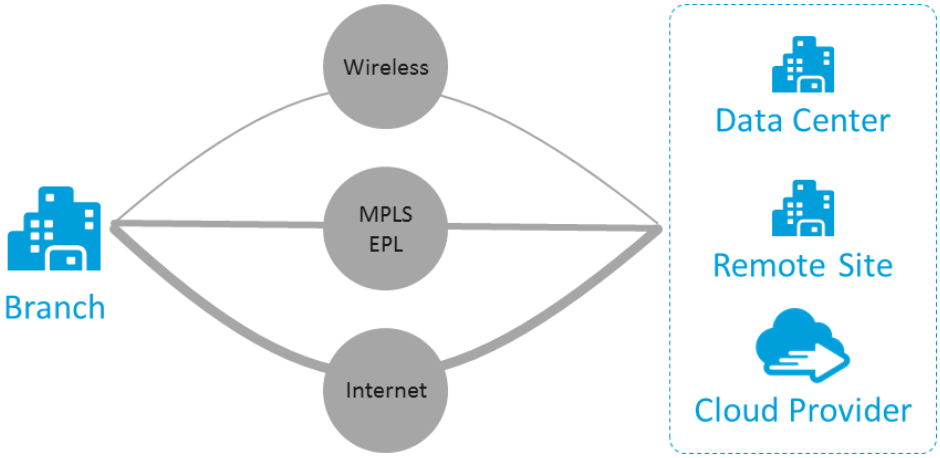

Figure 22: Different Types of Network Access & Transport

Dedicated versus shared access is one big question that has to be addressed. Dedicated access means that no other enterprise or service can use the access bandwidth. Examples include T1, SONET, and private microwave. Shared access allows multiple entities or services to share the underlying access. Examples of shared access include Ethernet, Cable Internet, DSL, and LTE. Shared access is cheaper but poses the challenge of guaranteeing bandwidth and associated performance along with security.

60% of WAN costs and 95% of problems reside in the last mile of network access.

The next access question is based on the amount of bandwidth required at the site. Fiber is by far the highest speed and best solution. Where it is not available, when it is too costly, or a redundant path is required, then copper on the wireline side can be used, and wireless options include 4/5G cellular, microwave, or satellite. Fiber access has the advantage of having nearly unlimited capacity and the best Mean Time Between Failure (MTBF). Copper can be installed the quickest and has a better Mean Time to Repair (MTTR).

In any WAN refresh project, the network access strategy is the most critical in determining total cost of ownership, agility, performance, and reliability. There are databases that track who provides fiber, copper, and wireless services to which business addresses or services that are available in the area. There is no single network service provider that has access everywhere. The following decisions must be made on deciding who manages the network access:

1) **Network Service Provider** – AT&T, Verizon, and CenturyLink all have a good footprint in their legacy markets, but outside of these markets, they have to buy access from someone else. The Network to Network Interface (NNI) between carriers used to be primarily layer 3 based (IP or MPLS) especially internationally. These days, most of the access NNI are layer 2, Ethernet-based. One of the challenges is that an NSP may not have connectivity to every access provider, in every market.
    i. **Pros** – Single or dual carriers and contracts with SLAs. The NSP manages everything with the option of the customer edge router (which is ~80% common globally, but only ~35% in U.S).
    ii. **Cons** – Multi-year contracts can limit flexibility, provisioning of faster/new access can take months, may not be the lowest cost.
2) **Virtual Network Operator** – Mettel, Earthlink/Windstream, Global Capacity/GTT which buy wholesale access from NSPs, ISPs, and fiber providers. While a VNO may have their own backbone, they rely on the best local access provider for last mile connectivity. Many times, VNOs can get access cheaper from an NSP since they buy it wholesale, rather than a large enterprise that has a big discount off of retail access rates. The VNO manages all the different access providers — both wireline and wireless/cellular — and they provide an enterprise with a single entity for ordering,

provisioning, support, and billing. Many of the VNOs have a heritage of delivery Plain Old Telephone Service (POTS) lines and ISP services since a large enterprise did not want to work with all the regional providers. Thanks in part to the 1996 Telecom Act, legacy NSPs had to open up their last mile networks at wholesale rates.

i. **Pros** – Lowest cost to a lot of small sites that are highly distributed, still a single company and contract to work with.

ii. **Cons** – Do not have the name recognition of the larger NSPs and are primarily focused on the North American market.

3) **Enterprise** – Large enterprises have large teams to manage many different NSPs or they will go out and buy or build their own fiber networks. This is especially common in some verticals such as manufacturing where a company primarily has just very large sites that need 10+Gbps of connectivity.

It is recommended that enterprises create network templates based on the size and importance of their site. The typical breakout is listed in table 3 below:

#	Site Type	Examples	Reliability	Bandwidth	Access Strategy	Transport Strategy
1	Core	Data Center, Cloud	99.9999%	100+ Gbps	2+ diverse fiber paths to Co-lo	EPL, VPLS, MPLS, Internet
2	Large & Mission Critical	Hospital, Manufacturing, Campus	99.999%	10+Gbps	2 diverse fiber paths or microwave/Pr-5G	EPL, VPLS, MPLS, Internet
3	Critical	Store, Bank Branch	99.999%	10-1,000Mbps	Fiber and copper and/or wireless	MPLS, Internet
4	Important	10-100 People in office	99.99%	10-100Mbps	Fiber and copper or wireless	Internet
5	Basic	1-10 People in office	99.9%	10Mbps	Copper or Wireless	Internet

Table 3: Site Type and Associated Access and Transport Strategy

Over time, fiber will become predominate and copper will fade away on the wireline access side, with its ability to carry 192, 100Gbps Ethernet connections on a single fiber. 5G cellular (both private and public) will become predominate on the wireless side with the capacity to support hundreds of 100Mbps connections within a couple mile radius. Fiber is subject to physical cuts that take a long time to repair, and wireless is subject to interference due to weather (heavy rain, snow, or fog) or congestion, so having diverse fiber connections or a combination of access methodologies is the only way to ensure high reliability.

Transport – The Network Underlay Paths

With the dedicated versus shared access strategy determined, the next decision is the OSI layer 2 (frames) and Layer 3 (packet) strategy. This takes care of getting the network traffic from the local NSP central office or co-lo to other NSP central offices and/or co-locations.

On the wireline side, Ethernet is the common and standard way of providing new layer 2 connections. Ethernet can be ordered in increments of 10 from 10Mbps up to 100Gbps. Ethernet connections can be point-to-point and ordered as an Ethernet Private Line (EPL) or a multi-point service called Virtual Private LAN Service (VPLS) than can span hundreds of sites.

Then the next critical question is whether to use private networks such as an NSPs MPLS network, or the public Internet. Many enterprises report that their contracted bandwidth rates for MPLS and Internet are the same, when using dedicated access such as 1Gbps Ethernet connection running on top of fiber. Where Internet is significantly cheaper is when there is a shared access model such as DSL, Cable, or Ethernet.

In the past, enterprises had to pick one or two network service providers and a single transport technology in order to provide the site to site routing, Quality of Service (QoS), and domain of their network. The beauty of SD-WANs is that these functions have been moved up into the overlay and so the underlying transport can be a combinate of public and private connections from a slew of different service providers.

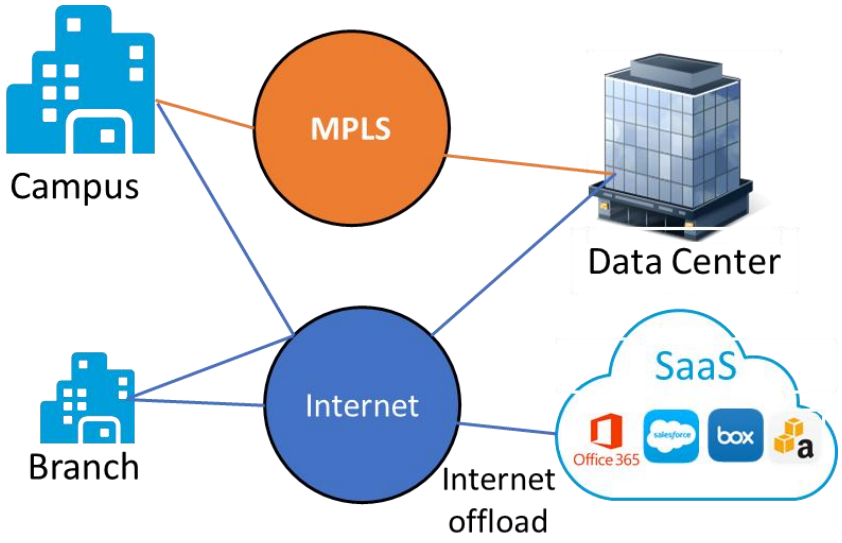

Figure 23: Hybrid Network Connectivity

Thus, the decision on whether to use EPL, VPLS, MPLS, or Internet to a site comes down to what the business requirements are. Many enterprises choose a combination or a hybrid approach. So, an enterprise has three choices in the underlay of connecting to a cloud provider.

1. **Internet** – An IPsec tunnel from the enterprise to the cloud provider
2. **Ethernet Interconnect** – An ethernet connection from the enterprise to the cloud provider
3. **MPLS Direct WAN** – Leverage the existing MPLS network and add the cloud provider

Each of these methods have their pros and cons as shown below in Figure 24.

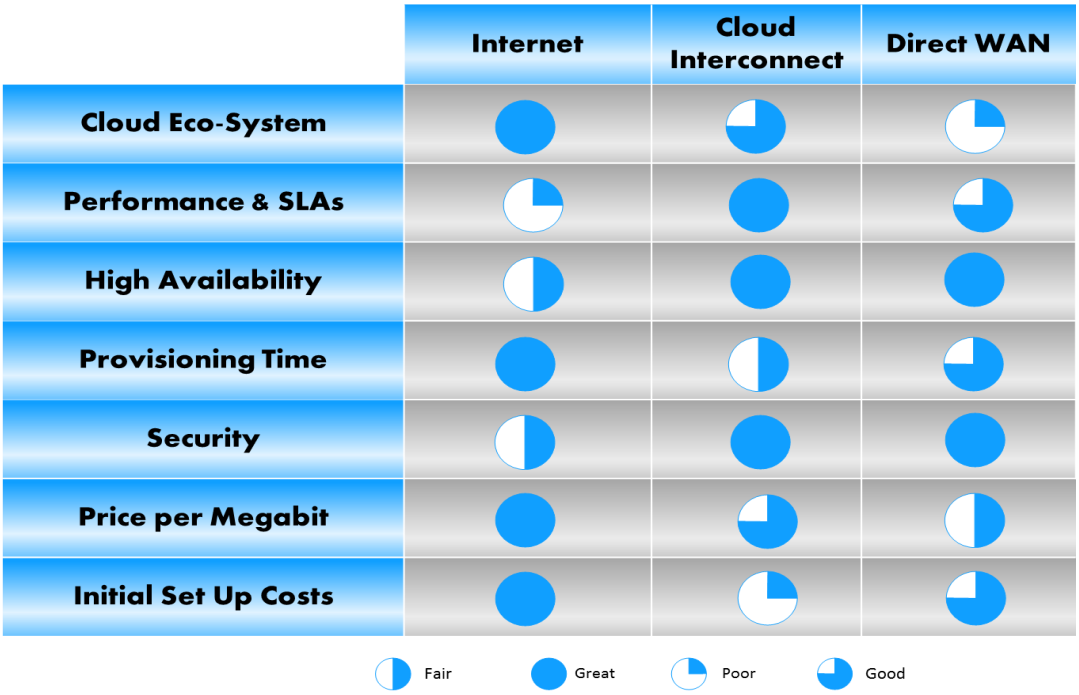

Figure 24: Pros and Cons of Network Transport Options

One thing that SD-WAN offers is the ability to use multiple connectivity types. So, if an enterprise wants to run a non-mission critical application in the cloud or test a new application, then Internet connectivity may be good enough. As scale, latency, and security requirements increase then cloud interconnects and/or direct WAN connections come to play.

Choosing the Right Hardware

One of the promises of SD-WAN is being hardware agnostic and not dependent on a black box. Additional network functions can also be run on this commodity off the shelf (COTS) hardware. The Network Function Virtualization (NFV) enables multiple network functions to run on a generic underlying hardware, instead of buying proprietary hardware for each network function and stacking them on top of each other.

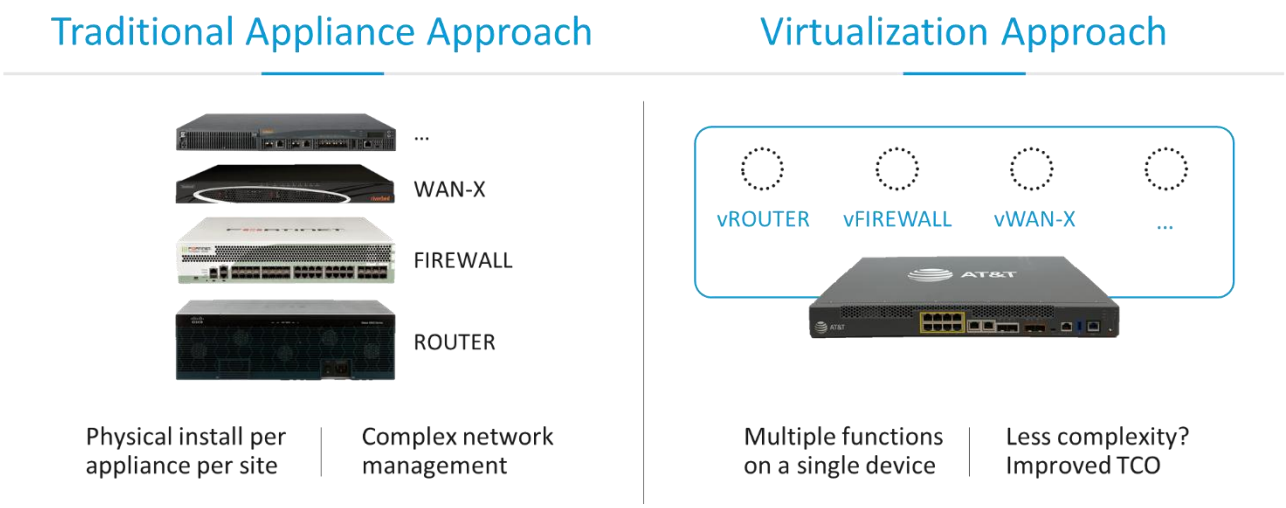

Figure 25: Example of a Virtualized Approach to Network Functions

Virtual Customer Premise Equipment (vCPE) provides network routing, security, and other functions via software versus dedicated hardware devices. While most enterprises have a virtualization strategy for their data center, few have one for the branch office.

Ask any CIO what their virtualization strategy is for the branch office, and they will tell you that all applications are moving to the cloud (private, public, or hybrid), and that they do not need one. Ask them what about routing, security, wireless LAN, printing, and other services that still reside in the local office, and their answer will still be the cloud. The reality is that some functions still need to remain local.

Redundancy is another large consideration in the branch office including if redundant hardware is required. Having a single point of failure is not acceptable for those sites that need reliability above 99.9%. There are also different HA redundancy models such as active/standby or active/active and if one device and/or circuit fails, all the traffic moves to the available device or circuit.

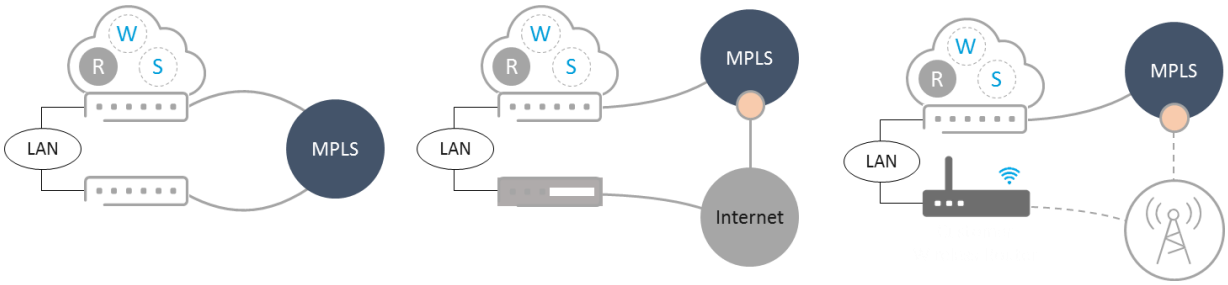

Figure 26: Hardware and Network Redundancy Examples for the Branch Office

When sizing the hardware used for a site, the following variables come into play for the network portion:

1. **Bandwidth** – Peak throughput, bi-directional, all interfaces
2. **Number of Tunnels** – Architecture of mesh vs hub and spoke and number of hubs including cloud gateways.
3. **Level of Encryption** – 128 vs. 256, hash algorithm, adaptive (recognizes if TLS is used at the application layer and does not re-encrypt at network layer).
4. **Use of FEC** – None vs. packet duplication vs adaptive (adds FEC as packet loss goes up dynamically)
5. **HA Model** – All traffic goes through one router (active/standby) vs traffic and links are split between both site routers (active/active), and flow/session state management
6. **QoS Levels** – Basic vs. complex such as going over an LTE link where buffering and retransmission are common, and bandwidth is asymmetrical and limited.
7. **Packet Size** – Traffic mix - Lots of small packets for voice vs. lots of large packets in file transfers
8. **Fragmentation** – Percent of packets that are being fragmented and reassembled
9. **Management Data** – Amount of management data being collected along with the duration of it being stored locally, for instance adding a time stamp and sequence number to every packet and measuring the inter-packet latency, jitter, and packet loss.
10. **Growth** – How long is the hardware expected to last and what is the anticipated growth during that time. Many enterprises have a 3 or 4-year amortization schedules on their hardware, and typical enterprise WAN growth is 30% per year.
11. **Interfaces** – How many ethernet interfaces are required for HA, WAN and LAN, Proxy, and WAN Op. Also, USB for console and out of band access, video port, and legacy interfaces such asT1/E1.

For additional functions such as firewall, WAN optimization, SBC…these all consume additional CPU and memory. Most base level SD-WAN solutions require 4 core CPU and 16G or RAM for 200Mbps of encrypted throughput. To size a physical box, virtualized, or cloud solution, most SD-WAN vendors have recommendations based on bandwidth and number of tunnels as a baseline.

> The SD-WAN overlay is the "secret sauce" and every vendor's implementation of their overlay is different.

Data centers, communications hubs, and cloud gateways will need significantly more computing power than at a single site to terminate all the tunnels and handle the aggregated bandwidth. One of the advantages of moving to software is that hardware can be scaled horizontally. Instead of having a single router support 1Tbps with a standby router (active/standby), 10 routers that each support 100Gbps can be clustered together with the load shared across them. An 11th router can

be added for redundancy in case one of the ten routers fail and there is a need to continue to support 1Tbps. This is a more cost-effective way to support high speeds.

The SD-WAN Intelligent Network Overlay

SD-WANs add a layer on top of the traditional WAN stack. This is the overlay layer which is the additional capability added through adding tunnels, labels, and/or session state. Through use of the overlay and a more intelligent control plane, the overlay provides application level networking and security over and above the network level which focused on the path/link networking and security. The SD-WAN overlay is the "secret sauce" and every vendor's implementation of their overlay is different.

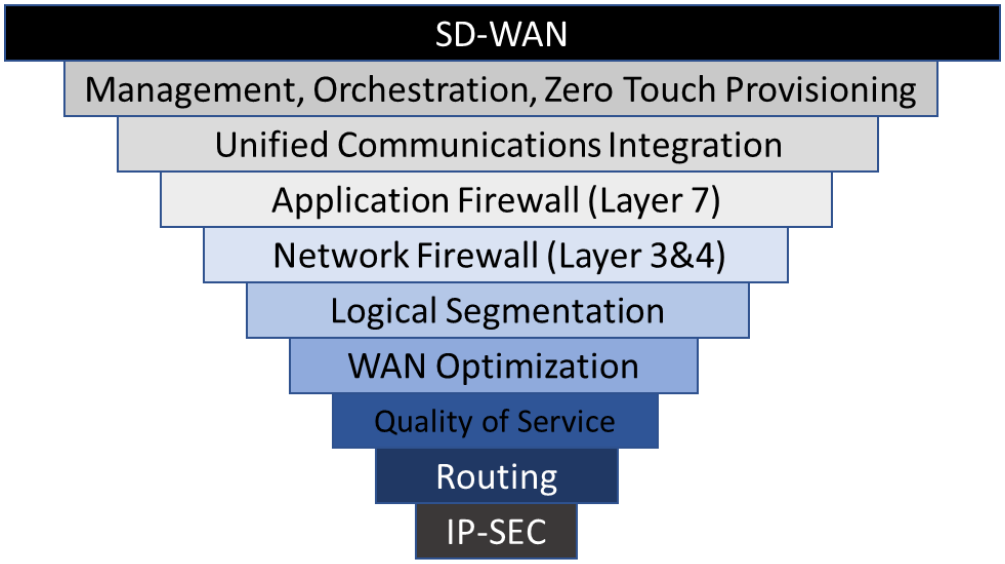

Figure 27: SD-WAN Network Stack

IPsec is the foundation of SD-WANs and provides the ability to route a packet via a path that the original IP address would or cannot follow along with providing the encryption. IPsec has many limitations, so SD-WAN vendors created additional functions to augment IPsec.

Every SD-WAN solution can be implemented in multiple ways, each with its pros and cons. With the business objectives and guiding principles set, the technical design comes into play. One design consideration is the connectivity of the IPsec tunnels in a point-to-point configuration from the site to a data center, or a mesh where every site has a tunnel to every other site.

Early SD-WAN implementations focused on a hub such as a data center, co-location, or cloud center that each branch office as a spoke directly connected to. In the hub and spoke model, a branch office must backhaul through a hub to talk with another site. While the underlay network may support an any-to-any architecture, the SD-WAN tunnels and their associated management are typically point-to-point. The latest SD-WAN architecture are a hybrid model where a branch

office connects both to the enterprises data center and to a secure cloud gateway so that all cloud and Internet bound traffic does not have to hairpin through the data center.

Routing

Overlay routing provides application flow control and per packet steering. The underlay network routes based on the shortest path and if the link is up. The underlay does not take into account the latency, jitter, or dropped packets of a link. Think of the SD-WAN overlay routing as Waze (Google GPS navigation) used in getting a car to a destination based on the fastest route, versus the shortest path or always using an interstate. So, if in the middle of an application session, the current network path falls below a threshold in performance, the network can intelligently route the traffic to a better path. Today's internet has many more brown-outs than black-outs, so the SD-WAN overlay gives one the ability to route around the brownouts while the underlay routes around black-outs. BGP and OSPF can be used to pass routing information between the underlay and overlay routing layer.

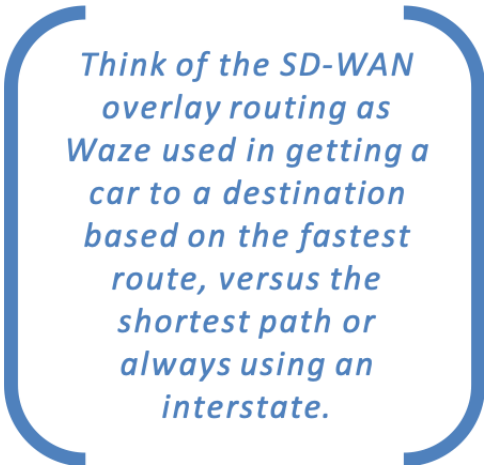

One of the challenges to underlay routing in world were mobile users are accessing cloud applications is that they are on different networks with multiple firewall/NATs in place. The underlay network only has visibility to the traffic within a domain or a single private network, not across many different networks. SD-WAN overlay routing can provide the end-to-end network performance controls and security. Overlay routing integration into underlay routing follows the standards of OSPF, BGP4/6, ECMP, and route reflectors.

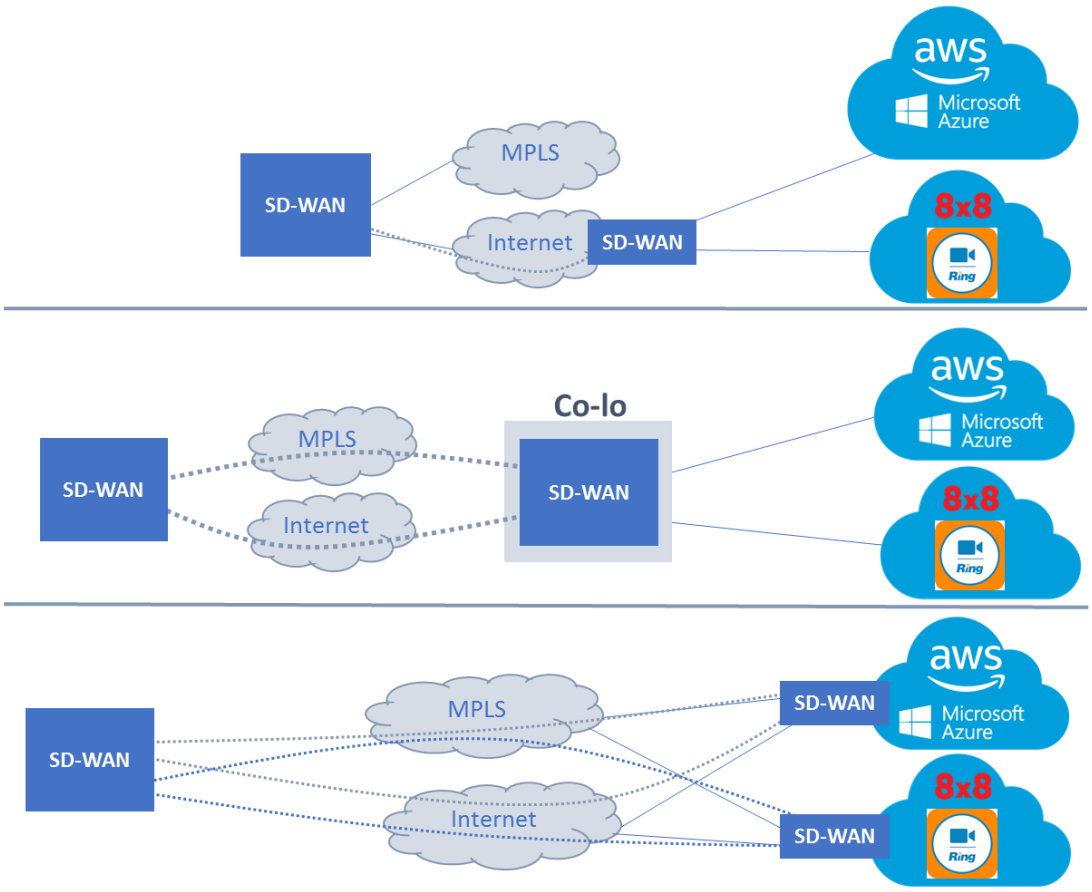

Figure 28: Three Options for SD-WAN Termination Location for Cloud Connectivity

Figure 28 shows where SD-WANs can be deployed. Many enterprises used SD-WAN initially as an access strategy when using the Internet as the access and transport strategy as shown in the top graphic. They then added their MPLS links into a co-location facility where they built communication hubs to internetwork their private and public networks and to get closer to cloud services. Today, leading enterprises are deploying SD-WAN all the way into cloud service providers to ensure end-to-end performance and security of cloud services.

Quality of Service

Overlay QoS can automatically identify and classify users, devices, services, applications, and data. While the underlay network uses IP addresses and TCP/UDP port numbers to identify and classify traffic, the overlay QoS is more intelligent and granular.

Different SD-WAN vendors have different strategies for providing QoS. Some vendors add sequence numbers and time stamps to every packet in order to track the latency, jitter, and

dropped packets on every flow/session. This applies to both TCP and UDP based applications but is done at OSI layer 3 in the SD-WAN versus the higher TCP or UDP/RTP layer.

QoS can be broken down into application visibility and control, traffic steering, and analytics as shown in Table 4 below.

Application Visibility	Traffic Steering	Analytics
Identification of applications	Route based traffic steering	Reliability and Availability
Application groups	App based path selection	Bandwidth usage
Application filters	Traffic Load Balancer	Bandwidth per application
Application logs	URL traffic management	Latency, Jitter, Packet Loss
App SLA enforcement	Voice Codec & MOS Score	Custom Report
App QoS – Traffic shaping		Exporting of data
App QoS – Rate limiting		

Table 4: QoS Controls and Reporting

Passing and utilizing underlay QoS using DSCP markings into the overlay and having the two work in concert together is tricky. Some service providers on their MPLS networks will honor the DSCP classifications whereas on Internet and Ethernet transports, this is ignored and all the QoS controls are done only in the overlay.

WAN Optimization

WAN Optimization (WAN Op) can be like security as a separate Network Function Virtualization (NFV) that is a separate software stack running in conjunction with SD-WAN on the same hardware. Most SD-WAN solutions have some WAN Op built in.

The degree of WAN Op required depends on the latency between the users and their applications. International and satellite connected sites that have latencies greater than 25ms can benefit from WAN Op. Enterprises that encrypt their applications with TLS, have a lot of real-time voice and video, and who utilize thin clients (where the application runs in the data center and the network just passes screens) will see little value in WAN optimization.

Forward Error Correction (FEC) is another feature available in most SD-WAN solutions. As mentioned earlier the author is not a big fan of this feature. FEC is built into TCP, so doing it at the layer 3 network layer only makes sense on high latency connections. FEC for real-time voice is best left up to an adaptive codec such as Opus which is end-to-end. FEC for UDP only makes sense if an enterprise is using a legacy voice codec such as G.711 or especially G.729.

A reason that two of the leading WAN Op vendors, Silverpeak and Riverbed, pivoted into the SD-WAN market early is that the WAN Op market is shrinking. Enterprises should baseline their applications and test whether WAN Op makes a significant difference in the user experience.

Management and Orchestration

Most SD-WAN solutions offer cloud-based management with the goal of simplifying and automating management of the network. Everything is done through nice graphical user interface and the days of needing to have a deep understanding of command lines for router configuration and operations are over.

This starts Zero Touch Provisioning (ZTP). An SD-WAN appliance is pre-provisioned with a cloud-based address to download its configuration once it has been powered up and plugged into the local site network. This "plug-play" model allows for a low cost and quick way to turn up SD-WAN.

This sounds good in theory, but making it work smoothly can be problematic because:

- Placing the equipment locally can vary by site, especially if cellular is going to be used and finding the place in the building with the strongest signal
- Getting all the right cables connected to the right ports done by a non-technical person is tricky
- Providing Internet access so the SD-WAN appliance so it can phone home, whether this is done via Wi-Fi, Cellular, or a network connection

Once the equipment is up and talking with the cloud management system, the application discovery and classification is automated on the leading SD-WAN platforms. The solutions even promise a single click for service insertion. The network is visually displayed, and the network operator can understand both the underlying health of the network along with the applications running on top of it. Many enterprises find that the application visibility and control as one of the top benefits of running an SD-WAN.

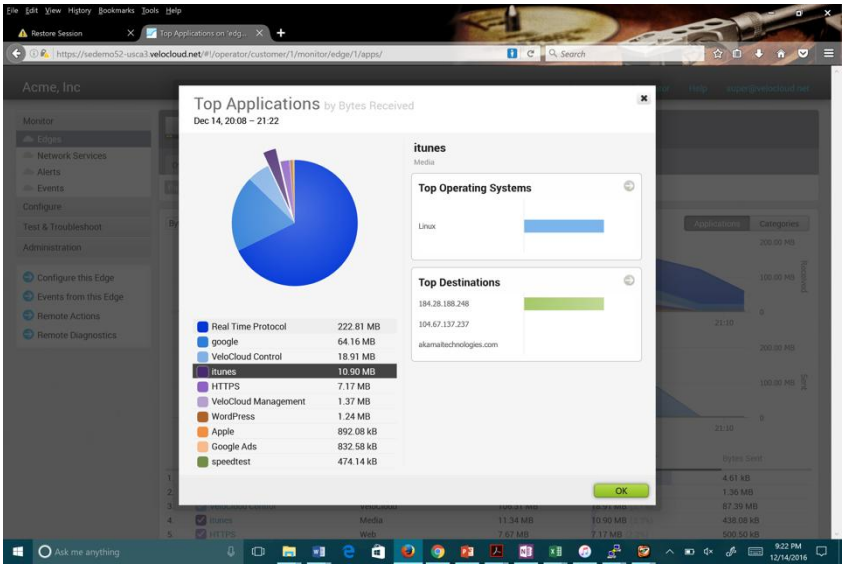

Figure 29: Example of SD-WAN Management Report – Top Applications

Security Architecture

Existing network approaches do not provide the levels of security and access control digital enterprises require. Security has become the top requirement for network designs, over and above the traditional requirement of reliability, performance, and costs.

Security is another network function that can be added to an SD-WAN as a separate software stack or utilize the SD-WAN security solution that is built. Most SD-WAN vendors provide layer 3 and 4 security rules based on IP addresses and TCP/UDP ports. More and more enterprises are moving their security stacks into the cloud and getting a managed security service.

SD-WAN SASE Strategy

There is a demand for immediate access for users, no matter where they are located or which device they are on, in a way that meets all security requirements. For too long, network routing and security have been separated. Merging them creates synergies over and above stacking the two together using Network Function Virtualization. Zero Trust Networking and Secure Access Service Edge (SASE) is gaining importance with many parties interested in how to incorporate these concepts into their SD-WAN solution.

> Zero Trust Networking at its core is the 1:1 micro-segmentation between these users, devices, services, applications, and data.

One of the advantages of SASE is modularity. This means using security components that are required instead of the entire security stack that comes with today's next generation firewalls. For instance, an enterprise can provide internet off-load at a campus location and create a policy that if the site is on the whitelist of approved sites and the application is TLS authenticated and encrypted with a validated certificate, then the user can route directly to the application such as Office365 or WebEx. All other traffic will then be directed to a more robust security stack that provides web filtering, sandboxing, DNS security, credential theft prevention, data loss prevention, and next-generation firewall policies.

Good security is intelligent, dynamic, and continuous. One-way SD-WAN and SASE empowers this is moving from security rules that are not just on the link, but one for every network session that connects users and devices to services, applications, and data. Zero Trust Networking at its core is the 1:1 micro-segmentation between these users, devices, services, applications, and data.

SD-WAN routing and security capabilities enable:

1. **Flexibility**: Using what is required versus a full, bloated software stack.
2. **Cost savings**: Instead of buying and managing multiple point products, utilizing a single platform dramatically reduces costs and IT resources.
3. **Increased security**: A Zero Trust approach to the cloud removes trust assumptions when users, devices, and applications connect. A SASE solution provides complete session protection, regardless of whether a user is on or off the corporate network.
4. **Data protection**: Implementing data protection policies within a SASE framework helps prevent unauthorized access and abuse of sensitive data.

SD-WANs are built on overlays such as IPsec in order to get an IP packet to route across a path that the native/original IP header cannot, along with providing path security via encryption. VxLAN is another overlay used by some SD-WAN vendors to provide segmentation and encapsulation over and above what one can do with a standard IP packet.

VxLAN offers a hierarchal, end-to-end method to segment network traffic to provide the performance and security controls that digital enterprises demand. While there is no overall SD-WAN protocol standard, VxLAN is an industry standard that can be used in data centers, cloud providers, campus, branch-office, and VPN solutions.

Why is VxLAN important? These are a few key benefits of VxLAN:

- **Scalability** – Traditional VLANs only scale to 4,096 unique networks within a domain. With Zero Trust requiring 1:1 micro-segmentation between users, devices, services, applications, and data, traditional VLANs do not scale. Security these days is not just about north/south segmentation, but east/west and micro-segmentation. VxLAN scales to 16 million unique networks within a domain.
- **Blending Virtual & Physical Networks** – The VXLAN VTEP can be implemented in both virtual and physical switches allowing the virtual network to map to physical resources and network services. VXLAN Tunnel End Points (VTEP) perform the encapsulation/de-encapsulation.

The “secret sauce” in utilizing VxLAN to provide a ZTN/SASE is mapping Identity and Access Management (IAM) directories to VxLAN. Directory Enabled Networking (DEN) has been around for decades, but has never taken off, in part because of scalability challenges, which is the same reason routers do not manage session states like firewalls do. But as networking and routing move to all software and platforms can scale horizontally, scalability is no longer an issue. Thus, the merging of identity and network perimeters.

SD-WAN Firewall

As enterprises and NSPs move away from a secured perimeter model to a Zero Trust model, security must start at the very edge of the network and maintained through the core. Deny-all is the

architecture that all IP networks will be forced to move to. With this, firewall functionality will exist everywhere within the network.

Firewalls can do the following things that a router cannot:

1) **Stateful** – Firewalls maintain session state for every TCP & UDP flow. This allows them to prevent denial of service and man-in-the-middle attacks as well as enforce security rules in both directions. When a packet is fragmented, the firewall understands the original size of the packet, but a router does not.
2) **Auto-Updates** – Security vendors and their products are designed to be updated in near real-time as new threats are identified. The number of hacks (both successful and unsuccessful) is growing by over 40% per year, and the methods for hacking are constantly evolving with a new one coming out every 12 seconds.
3) **Logging & Reporting** – While routers do provide logs, they were not designed from the ground up to do logging and logging severely impacts router performance. Firewalls are designed to log and correlate traffic. After a breech, having logs to understand exactly what happened is crucial.
4) **Least Privileged Access** – Firewalls are designed to block everything, and then specific rules are applied to allow traffic through. Routers are designed to pass traffic and have default routers, and rules are then applied to block traffic. The paradigm difference is important, both technically and culturally.
5) **Signatures** – Next generation firewalls have the ability to detect traffic patterns and bits within a packet and match it to known security threats.

With threats from exponentially increasing endpoints coming from IoT and BYOD which are commonly on networks outside of the core enterprise network, the levels of threats are also growing exponentially. Part of the reason the SD-WAN market is taking off is that the successful new SD-WAN vendors are integrating routing and security into a single platform. An SD-WAN security stack should have the capability shown in table 5.

√ Stateful Firewall	√ DoS	√ Filtering
√ Zone & Endpoint Segmentation	√ IP ICMP Flood prevention	√ Whitelists/Blacklists
√ Application, User, and Content Identification	√ TCP/UDP Syn Flood mitigation	√ Reputation
√ IDS/IPS & DLP inspection	√ Anomaly detection	√ History
√ Proxy – DNS, Web, Email	√ DDoS identification	√ Import from Service
√ Logging & SIEM integration	√ Threat Management, Analytics	

Table 5: SD-WAN Security Capabilities

SDN and NFV will not solve the problem of combining routing and security. Putting two separate platforms on a common hardware platform does not create the synergy of being able to use a single orchestration engine for all security and routing policies.

Segmentation

Segmenting networks using VLANs and VRFs will not meet tomorrow's business needs. Network segmentation logically separates traffic over the same physical network. Segmentation is used to isolate users and applications for security and performance requirements. The most common uses of segmentation are to prioritize real-time traffic such as voice and isolate credit card authorization traffic to meet PCI requirements.

The business needs for segmentation will be 1,000x greater in the future because of:

1. **Zero Trust Security** – Move away from parameter-based security to a zero-trust security model where no device on the network is trusted. Thus, every device and application are segmented from each other, and then there is a hierarchy to be able to manage groups. A true zero trust network uses whitelist routing, which only allows users and applications to get on the network if there is an explicit policy to allow it. No broadcast domains or default routes mean millions of segments in large networks.
2. **Edge Computing** – The digital world is where things (IoT) and users converge. Augmented reality applications are always on, providing a contextual, dynamic, and interactive experience that is hyper-latency sensitive. While centralized cloud data centers will be used for orchestration, the applications with AI and tons of local data will need to be widely distributed to thousands of distributed data centers. The internetworking between all the distributed data centers owned by hundreds of providers using many differently managed IPv4/6 networks, both wireline and wireless, adds millions more segments.
3. **Video** – Cisco forecasts that 82% of traffic on public and private IP networks will be video by 2020. There are different types of video applications that need to be segmented for both security and performance requirements. Figure 30 demonstrates nine different video segments based on levels of security trust and network performance requirements. Within each of these segments, further user and application security can be provided which creates hundreds of more segments.

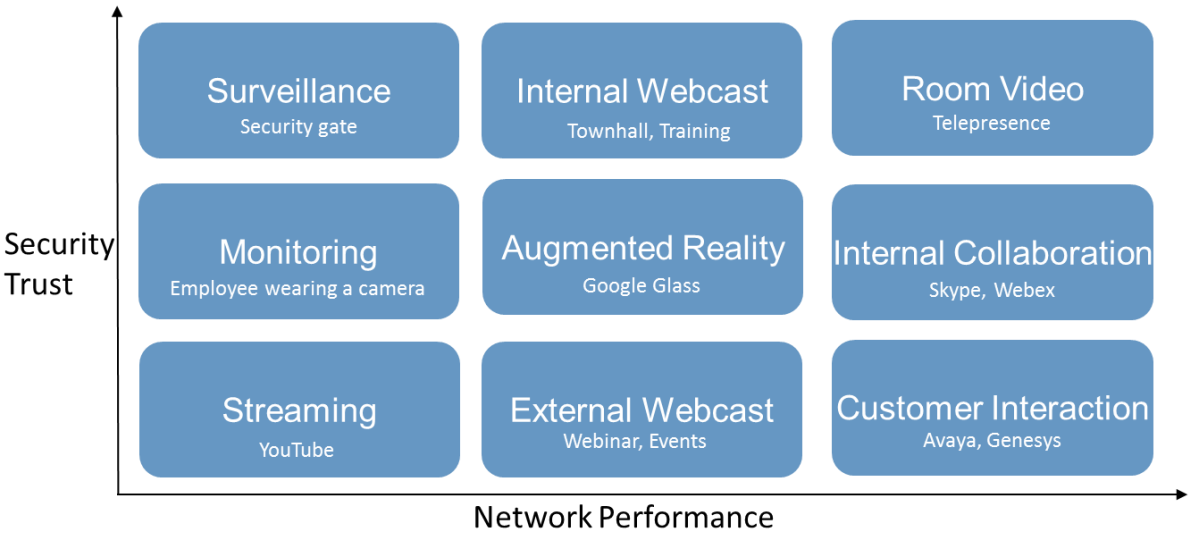

Figure 30: Video Traffic Segmentation Example

In order to scale network segmentation, network routing and security must work together, not diametrically opposed as they are with today's routers and firewalls. Application and endpoint identity and access controls using directories must be integrated with network routing and security policies.

The SD-WAN space is starting to provide this from the branch office to the data center and cloud, but to succeed SDNs need to go from endpoints of users and things all the way to the containers where the applications are hosted. One foundational challenge is that the segmentation technology in the LAN, WAN, Data Center, and cloud are all different.

> The segmentation technology in the LAN, WAN, Data Center, and cloud are all different.

If you are a retailer with 20 VLANs in a store, going across multiple networks (MPLS, Internet, LTE, VSAT) to multiple data centers and many different cloud providers, managing network segmentation through routers, firewalls, load balancers, and WAN optimizers leads to "ACL Hell." This complexity results in networks that are fragile, costly, slow to change, and unsecure.

Internet Offload or Not

One of the big design considerations in an SD-WAN architecture is where to provide access to the Internet. In legacy MPLS only WANs, all network traffic that was Internet bound routed through a data center and then through a full security stack before going out to the Internet. Many enterprises report that 80% of their network traffic is destined for the Internet, so backhauling traffic to a data center is both expensive and hurts performance by adding network latency.

The solution is to add Internet connectivity to every site and to offload the traffic directly. This requires firewalls at every site which adds additional costs and complexity. To get around this, many enterprises will create an IPsec or GRE tunnel to a managed cloud security service from providers such as zScaler. zScaler supports a maximum of 300Mbps for IPsec and 1Gbps for GRE throughput per tunnel. This can be limiting for large enterprises.

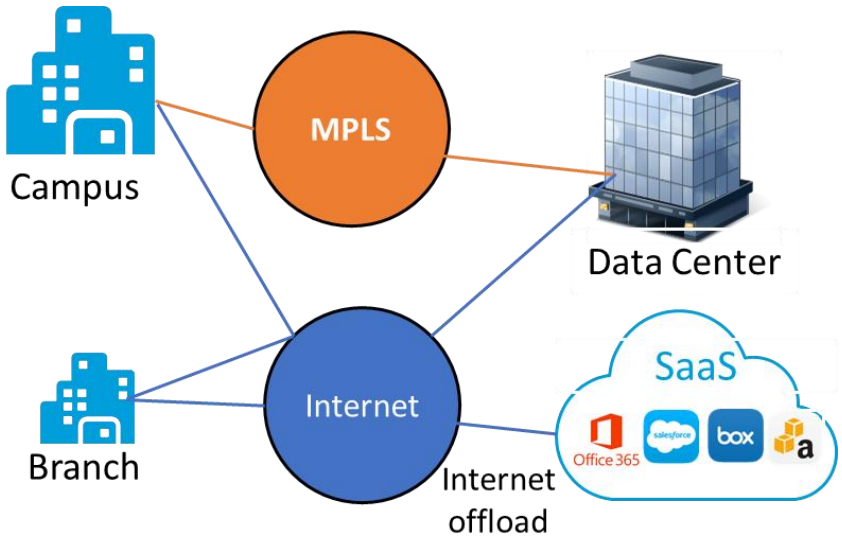

Figure 31: Internet Offload Example

Internet offload is for outbound traffic, with users going to Internet hosted services and applications. Inbound traffic for web hosting is a separate beast and should still be done in a data center (private or public cloud). The security required to protect hosted applications is very sophisticated and cost prohibitive to put at branch locations.

UC Architecture

Voice is one of the most critical business applications and is very sensitive to network performance. Voice is the "canary" on the network: if there are network performance problems, it is the first application to be impacted. Plus, because video consumes the bulk of network bandwidth, Unified Communications (UC) is being called out in its own section of something to architect.

Network Challenges with UC

Network managers have the responsibility of providing a high-performance, reliable, and secure network for the enterprise, but they have less and less control of what runs over it. Talk with any network manager about his or her top three problems — collaboration tools are at the top of the list. Adoption of freemium, cloud-based collaboration tools by an increasing number of user groups is exacerbating the "wave" effect on enterprise networks. This is because these tools use adaptive codecs for voice and video sessions.

The wave effect occurs when network utilization hits 100%, backs off, then hits 100% and backs off again in a continuous cycle. For example, TCP windowing starts the wave effect with a TCP session creating a large window and then backing off once a few packets are dropped. Having many flows across a single network link leads to large waves. Waves are problematic because every time network utilization hits 100%, applications are impacted, including those critical to running the enterprise.

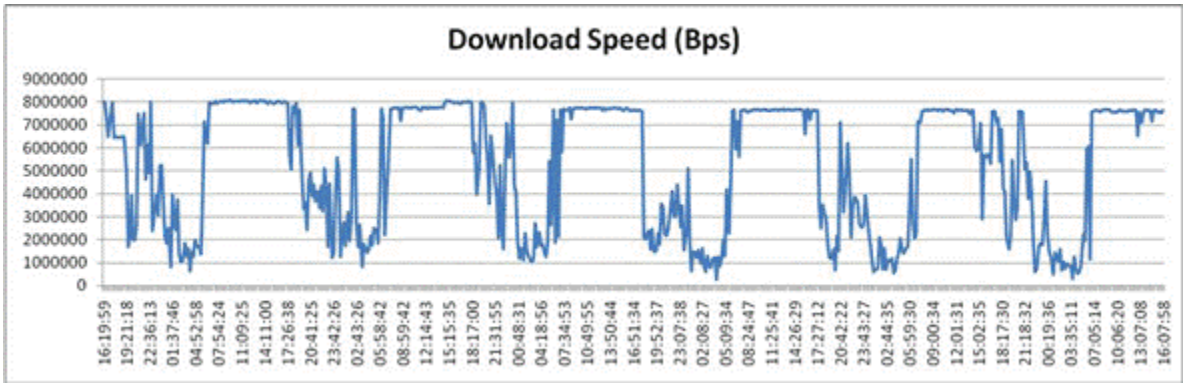

Figure 32: Wave Effect Caused by FEC or TCP Window Sizing

Adaptive codecs do the same thing but are worse because their back-off intervals take longer. An adaptive codec will first utilize Forward Error Correction (FEC) and send more packets before it backs off. When a 1-Mbps video conference experiences packet loss, for instance, the session will burst up to 3 Mbps for 10 - 15 seconds before backing down to 512 Kbps (by reducing resolution and/or frame rate).

Adaptive codecs work well when only a few sessions are running on a network. But get lots of sessions together on a single network pipe, and waves occur if the network pipe isn't big enough.

Many WAN optimizers do a good job managing TCP windowing, but struggle to manage real-time UDP optimization used by adaptive codecs. Compounding the problem, team collaboration tools use TLS or D-TLS encryption. Both make identification of the voice, video, and data traffic difficult within the encrypted tunnel.

Traditional network gear doesn't solve the adaptive codec wave problem. While call admission control (CAC) can limit the number of concurrent voice and video sessions on an internal collaboration system, it doesn't work on over-the-top cloud-based solutions. Unlike voice, which is a fairly regular stream, video has huge bursts of traffic.

Software Defined WAN players build in flow intelligence to be able to solve this problem in real time. They can:

1. **Send traffic across multiple paths** based on network packet loss, latency, and jitter.

2. **Prioritize voice above video** and rate limit the video as required, especially across networks that aren't honoring the DiffServ QoS settings, even within a WebRTC D-TLS stream.
3. **Identify and control real-time traffic**, even if it's coming from a cloud provider, based on flow characteristics — voice, for example, has packets of a fixed size that come at a fixed increment of time.
4. **Report mean opinion scores** (MOS) for each voice and video session for real-time alerting and long-term tracking.

One enterprise network manager said recently that users love Cisco's Webex Teams (formerly known as Spark), but their Cisco WAN can't manage the traffic effectively, causing intermittent voice on conference calls. It's okay for video and data traffic on a conference session to get delayed, but any interruption in the voice traffic has a direct impact on the QoE of the conference call. This is more critical than ever as the world deals with the COVID-19 pandemic and are replacing in-person meetings with conference meetings at massive scale.

Enterprises have three options for solving the adaptive codec wave problem caused by cloud-based collaboration tools:

1. **Block it** – Put strict firewall rules in place and try and block all cloud-based collaboration tools
2. **Overbuild** – Add big networking pipes so network congestion doesn't occur
3. **SD-WAN** – Dynamically manage the voice and video traffic riding the network

Option 3 is emerging as the most popular, since blocking traffic only makes IT more unpopular than it already tends to be and overbuilding leads to spending too much. The emerging SD-WAN vendors can solve this problem elegantly with new ways of bandwidth shaping across multiple paths, while monitoring the quality of every session. Traditional network vendors will reroute a voice/video session only if the link is down, while leading SD-WAN vendors can reroute traffic in less than a second to a better path if dropped packets or jitter exceeds a pre-defined threshold.

Cisco says that by 2022, 85% of all IP traffic (both business and consumer) will be IP video. Managing video and the waves of traffic caused by adaptive codecs is one of the drivers of the SD-WAN market.

Adding SBCs to SD-WAN

Session Border Controllers (SBCs) are another network function that can be bundled in with SD-WAN. Retailers have adopted SD-WANs in a big way and high quality, reliable voice services are imperative because of the hyper-competitive retail market.

Customers call retail outlets all the time to see if the latest hot toy is in stock, check pricing, see if their prescription is filled, and check store hours. If this call is not answered in a timely manner, the odds are high that a customer will go to a competitor. What is an even worse

experience is if the customer calls the store, and half-way through the conversation, they are disconnected.

Traditionally, retail stores would have a local PBX or auto-attendant and analog POTS line. This model is expensive, non-agile, an un-integrated. As retailers work on their store of the future, they want to centralize the telephony into the cloud (private or public) to reduce costs, integrate voice with other communication channels, and be more agile. Part of this strategy is also to move some users to softphones, while still providing a hard phone to the power users who support multiple calls concurrently. One challenge is that this migration takes time and during the transition there can still be legacy equipment and circuits within the store.

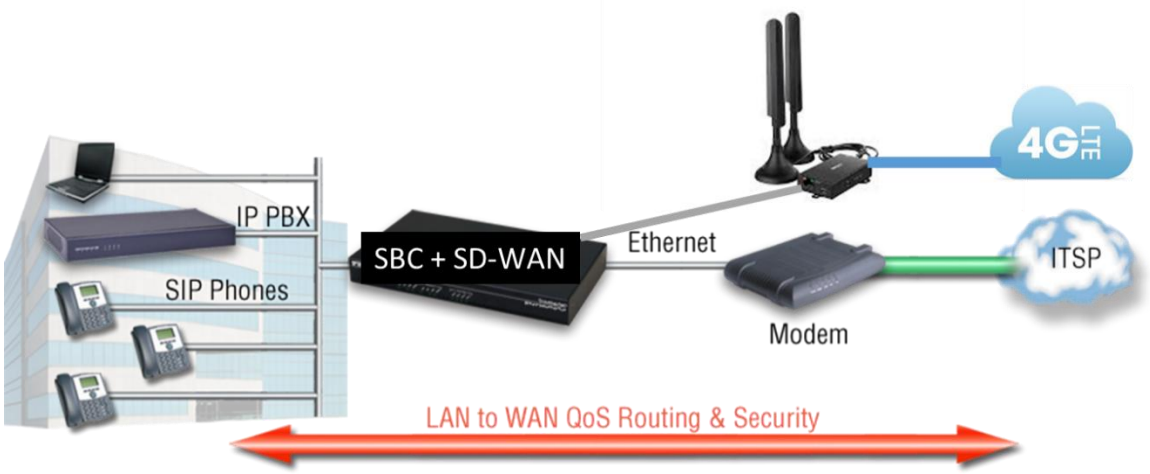

Figure 33: SBC and SD-WAN Together

The SBC provides:

1. **Analog interfaces** if required at the store for things like paging systems, walkie-talkies, old fax machines, and old phones.
2. **Voice Quality** with the ability to monitor voice quality and troubleshoot any issues that may arise.
3. **Centralized signaling** while keeping the media distributed.
4. **Voice Security** – SIP signaling and media security.

The SD-WAN provides:

1. **Network Quality** – The ability to reroute to a better network path when network congestion occurs in under 1 second, allowing enterprises to use Internet and LTE network links instead of MPLS.
2. **Network Security** – Network segmentation from store LAN, across the WAN, into the cloud and encrypting the traffic if required.

3. **Direct Routing** – Instead of bouncing through a data center, being able to route directly to a UC platform such as MS Teams, or Webex Teams.

Together, both can provide centralized orchestration and management to meet current and future requirements. Video calling (one or two-way) is part of the vision of many retailers. For example, an auto part retailer tested out video calling. Their representative could see the auto part their customer was looking for and offer advice. They were even successful on bringing in third-party video from YouTube and from another store where an employee had expertise in specific makes and models of cars. This is just one of many examples of how video calling will improve customer service.

Faster + Better + Cheaper + More Secure is the goal of all great solutions. In this day and age, why make trade-offs when you can have it all?

Testing VoIP in SD-WAN

Happiness is expectations minus reality. So, what should one expect when testing VoIP on a SD-WAN solution? The expectation in the branch is to get even better performance and security than MPLS, while also avoiding the high cost or complexity of MPLS. Since VoIP is the most sensitive application on most IP networks, getting it to work well will ensure the SD-WAN solution will be successful.

Here are some things to test with VoIP on an SD-WAN solution:

1. **Failover** – Using multiple links to ensure application performance. If one link fails or has a brown out where dropped packets or jitter goes above a certain threshold, the voice conversation should fail over to a better path in less than 2 seconds. This requires real-time link monitoring on a sub-second basis. Traditional IP routing protocols take 5 - 30 seconds, and in most cases only fail over if the link is down, not if UDP packets are being dropped. Failing back to the original path once it is healthy should also be tested. Especially test this with LTE, since most WAN outages are in the last mile.
2. **Load Testing** – If a branch office needs to be able to support 20 concurrent calls, for example, this would traditionally take 1,200Kbps, assuming a G.729 codec. In an SD-WAN world, depending on SD-WAN vendors overhead, this would take 2,720Kbps. Yes, SD-WAN overhead is significant! DSL, cable, and LTE provide asymmetric bandwidth meaning download speeds are significantly higher than upload speeds. While this works well for Web apps and downloads, voice is symmetric and requires this amount of bandwidth bi-directionally.

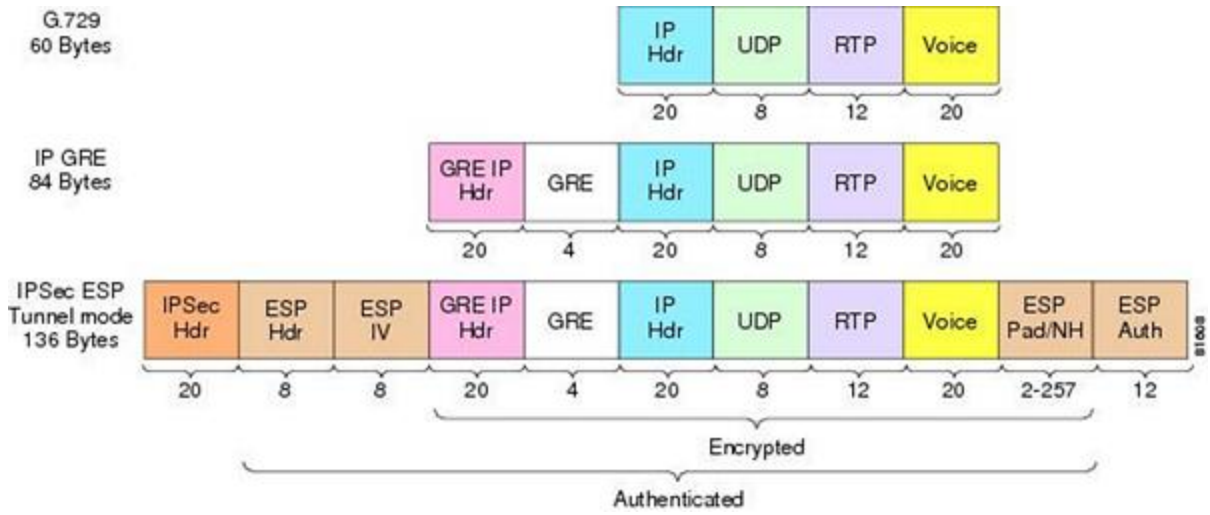

Figure 34: Example of SD-WAN Overhead

3. **Peer-to-Peer** – When calling from an office across town to the warehouse, the VoIP packets should not have to hairpin through a data center or remote cloud provider. Many early SD-WAN deployments are hub-and-spoke based and that adds latency. With large implementations, peer-to-peer has trouble scaling and it's like going back to the frame relay days. We moved to MPLS because of the peer-to-peer capability, and we do not want to move backwards.
4. **MOS Alerting & Reporting** – Measuring packet loss and jitter in *both* directions and creating real-time alerts and longer-term SLA reports is critical. As a network manager, when call quality is poor, you are guilty until proven innocent. So, having the data and being able to correlate it to specific calls is critical.

Future Architecture Considerations

SD-WAN technology and best practices are still evolving. An SD-WAN architecture is a blueprint that represents the best design at a point in time. As enterprises create their future SD-WAN strategies they should factor in the topics and recommendations we'll describe in this section starting with the role that open source can play.

Open Source's Role

While not making headlines or being overhyped like artificial intelligence (AI), open source software is fundamentally changing how we build our future networks and communication services.

Open source software such as Red Hat OS, WordPress for blogging, and the Google

Chrome browser (among others) has radically transformed various parts of the IT stack but has been more of a niche solution in networking communications. Some of the reasons for open source's lack of success in this space are:

- **High Availability** – Good enough is not good enough. Businesses cannot run if their network and communication systems are down.
- **No Jitter** – Real-time voice and video do not tolerate variable delay, which is why they were the last applications to be virtualized. Even today, many network and UC applications run on black or gray boxes.
- **Security** – The perception that open source is less secure than commercial software persists.
- **Complexity** – Router and PBX software has millions of lines of code and thousands of features. Look how long it took Cisco to make an IP phone system that could compete against a traditional PBX!

This is all changing as cloud providers and Digital Enterprises adopt open source-first strategies, and UC and networking becomes truly software based and free from the shackles of proprietary hardware. In an all-software world, communications get embedded directly into applications, so you don't have to leave your business application to place a call or launch a video chat. Network switches and routers, along with the communications applications, can be highly distributed without a single point of failure. This helps overcome the challenges of traditional monolithic, proprietary, and expensive solutions.

So, open source's impact on the network markets will be:

1) **Additional Vendors** – The barrier to entry will drop as new vendors and cloud service providers leverage open source for as much as 90% of their offerings, and then add their value-added capabilities on top.

2) **Lower Margins** – Even enterprises that are laggards in technology adoption can leverage the open source ecosystem as a bargaining chip.

3) **Faster Evolution** – An all-software market means we'll see providers adding features more quickly than ever to support the always-on-and-interacting digital world.

4) **Better Security** – Every week, there seem to be security patches for closing backdoors or vulnerabilities in today's proprietary solutions. An open source and collaboration community identifies and fixes problems more quickly than a single vendor does.

One example of a recent networking-related open source initiative came in 2019 from flexiWAN which introduced an open architecture and open source SD-WAN solution. The 60 or so SD-WAN vendors in the market today already utilize a lot of open source in their products, thus going completely open source is the natural evolution. SaaS providers in particular will find value in being able to embed SD-WAN into their offerings to provide end-to-end performance and security of their applications.

Standards groups will need to evolve too, as the focus of work changes from how technologies integrate to orchestration of a common software stack. Perhaps it's time to update coding skills and join the software revolution that's quietly empowering more open source solutions in networking and communications markets.

IPv6

Enterprises should limit their investment in IPv6 and see it as a tactical protocol. IPv6 is 25 years old and was a tactical protocol designed to correct the limitations of IPv4, such as the limited number of IP addresses. As enterprises exhaust their IPv4 networks, they should look past IPv6 as the answer to their network challenges. IPv6 has the following limitations:

Enterprises should limit their investment in IPv6 and see it as a tactical protocol.

1) **Additional Overhead** – For voice traffic, IoT, and other network traffic that has small packets, the additional overhead associated with IPv6 increases bandwidth requirements.
2) **Minimal Internetworking QoS** – Going from IPv4 to IPv6 and between NAT boundaries causes QoS and routing policy information to be lost.
3) **Zero Application Awareness** – This is still a packet-by-packet routing technology and is unaware of flow and stateful session security.
4) **Lack of Security** – The Internet Architecture Board recommends that all future protocols support end-to-end encryption.

The killer application for IPv6 is larger address space but at a price of additional overhead. Yes, there are some other advantages to IPv6 including multi-casting, but nothing that is driving enterprises. With only ~12% of IPv4 addresses being used, the IPv4 address problem is more of an allocation challenge than lack of addresses. We can limp along in the IPv4 world for another decade until something truly better comes along. Named Data Networking (NDN) has the potential to be the truly better approach as the type of traffic on the Internet changes (lots of video and content), and the needs of the Internet change (more security).

Eight Networking Disruptors

IP Networking has not changed since the mid-90's. For the first time in 25 years as software defined networking takes off, disruption is occurring in the networking field. SD-WANs are a start, but they are more of an evolution of IPsec tunneling versus a full disruption. Below is a list of significant potential networking disruptors coming in the next five years.

1. **Private 5G replacing Wi-Fi X** – For new construction of a large building or campus, 5G is faster, better, cheaper, and more secure than Wi-Fi 6. For companies that run PoE and 100/1000Mbps to their access points, 5G vs Wi-Fi 6 will need a careful analysis since Wi-Fi 6 & 7 are best served by 10G fiber. For companies that just need more wireless

capacity and have existing Wi-Fi infrastructure, Wi-Fi 6 makes sense. *Prediction: 5G will move into the WLAN market.*

2. **Using Multiple Networks Concurrently** – From a mobile device, being able to use multiple 4/5G and Wi-Fi networks concurrently instead of having to roam from one to another. eSIMs allow one to use multi-cellular providers. For power mobile users the ability to use multiple networks concurrently to increase bandwidth, lower costs, and provide another level of security is worth it. *Prediction: Faster, Better, Cheaper, and More Secure wins in the marketplace and power consumers will demand the real-time freedom of using whomever they want, when they want.*
3. **Zero Trust Networking** – Zero Trust is gaining momentum and was one of the top words used at RSA this year. Zero Trust Networking (ZTN) is where no packet gets onto the network without prior authentication and authorization. Thus, all the malware and malicious traffic is stopped at the edge of the network versus in the middle, or on your virtual doorstep. *Prediction: ZTN is the only way to truly secure networks and networking security will become the top priority of next generation networks over and above being faster, better, or cheaper.*
4. **Doing away with IP addressing** – IP addresses themselves are not going away anytime soon, but their importance will as SDN takes off, and the control and management planes use words to define routing and security policies and integrate with DevOps tools. *Prediction: IPv6 will continue to be of little value to any organization and future routing and security protocols will be based on words as software eats the world.*
5. **Opensource** – Opensource has disrupted every area of IT with networking becoming one of the last bastions for opensource to become widely adopted. The big cloud companies all built their own networking solution because there was nothing on the market that could meet their needs cost effectively. Plus, in a software world, routers can scale horizontally, so the days of paying $1 million for a router that handles 1Tbps of throughput are over. *Prediction: Opensource is coming. Embrace it and learn how to code.*
6. **Cloud Networking into the Enterprise** – Networking and security are the top barriers for cloud adoption. In order for a cloud company to effectively deliver end-to-end quality and security, they must manage the network between their applications to their users. Microsoft is one example of a company embracing this strategy with the rollout of their Virtual WAN. Equinix with their virtual co-location is driving networking closer to the edge of companies versus making companies backhaul their traffic to just a few co-locations. Many UCaaS providers have adopted SD-WAN to help ensure quality. *Prediction: Networking is the bane of any cloud company, but more and more will get into or partner with someone to deliver their service directly into the enterprise.*
7. **End-Point Multi-Cast** – The killer application for the Internet is "Look at me now." Facebook, Instagram, and YouTube are all built for the user to send the information to

them, and they disseminate it while they monetize your privacy and content. True freedom is for a user to be able to directly send all their content directly to their followers. This will require multi-cast directly from the endpoint. *Prediction: Facebook will be replaced/augmented over time by distributed, more private user-to-user services.*

8. **5ms Network Latency for Apps** – VR/AR, auto-driving, and next generation applications will require that users are within 5ms of the applications that they are consuming. No more VPNs and network backhauling to a data center. Networks will be redesigned to connect users via the best path, not what BGP thinks is the shortest path. *Prediction: The trend of Internet exchange points being more distributed will continue and eventually there will be one in your neighborhood.*

SD-WAN Predictions for 2020

1. **Acquisitions** will happen because while the SD-WAN market is growing rapidly, it is not big enough to support the 60 SD-WAN vendors. There will be consolidation and mergers to gain size. The large cloud vendors are getting into networking and acquiring their own SD-WAN solution(s). Go big or go home!
2. **Security and SASE/ZTN** strategy becomes the biggest driver for vendor selection. Part of the appeal of SASE is modularity. This is different than NFV. End-to-end, 1:1 micro-segmentation strategies are in vogue.
3. **Open source** plays a bigger role with major NSP backing such as AT&T and Telefonica. More vendors, especially security, get into the SD-WAN space and the number of SD-WAN vendors grows from ~65 to 90 even with some acquisitions as noted above. Software pricing normalizes to around $10/month per site or instance.
4. **Continued Revolution vs. Evolution** as enterprises either go all in or continue to do POCs and dabble with the technology, many hoping that Cisco/Viptela will someday be simple and run on their existing ISRs.
5. **Multi-vendor SD-WAN strategies** become common among enterprises and service providers depending on use case (Small Branch, Campus, Partner, Cloud IaaS, Cloud SaaS, DCI, VPN/user). One vendor/solution to rule them all is a pipe dream that is held by those who have used Cisco for decades.
6. **Interoperability will become more important**, first at the management/orchestration layer and then over time at the protocol level. *Who will create the SBC for SD-WAN?*
7. **IaaS & SaaS** get more serious about networking as the demarcation point moves closer or even into the enterprise for cloud services. MS-vWAN is an example and this will disrupt the legacy telco business model. The fact is that the large cloud providers have bigger backbones than the Telco's. Edge computing is an example of cloud getting closer to the enterprise and users. Telco is only the last mile with technologies such as 5G.

8. **DevOps Is the New Buyer** as software eats the world down to the networking layer. Elastic, dynamic, integrated networks are required. DevOps does not care about how the plumbing or electricity utilities work they just want to consume them on-demand.
9. **AI, ML, IBN** will continue to be over-hyped. Every network is unique, and things happen that cannot be predicted. Tightly managing networks will continue to be elusive which is why many cloud providers will continue to overbuild their networks.
10. **Hardware headaches** will continue, for as much as networking is moving to software, the hardware still matters to ensure optimal costs, performance, and agility.

Deploying an SD-WAN

The process of starting an SD-WAN to the deployment of a few sites can be done in four months, or 16 weeks. Figure 35 below shows the primary steps of gathering the business requirements, translating these into an architecture and then in turn a lower level design.

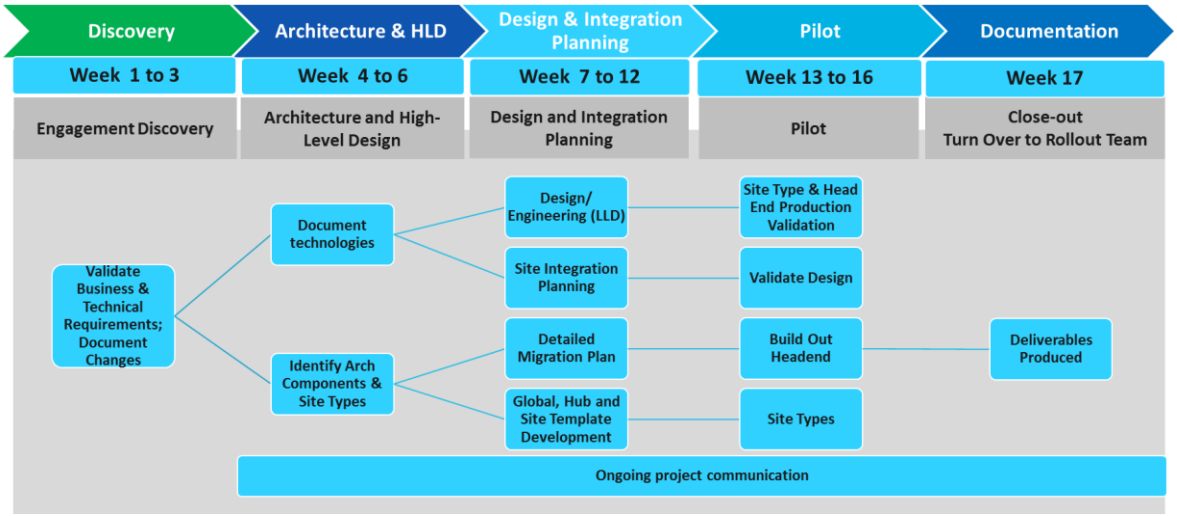

Figure 35: Example of a 16-Week Plan to Architect and Pilot SD-WAN

In deploying an SD-WAN there are numerous logistics to keep track of, which are covered in this section.

Guidelines

- Roles & Responsibility Matrix
- Master System of Record
- Equipment Location & Physical Logistics
- Order validation for all transport services

There are several factors to consider when deploying a solution that requires installation of physical infrastructure at a customer location. Both physical and logical implementation

instructions must be precise as it becomes costly in terms of time and money to correct installation errors. The following best practices are recommended to ensure a successful installation.

Roles & Responsibilities

A Roles & Responsibilities Matrix with SLAs should be created and agreed upon by all parties to ensure expectations are inline.

- Key Organization Contacts and Escalation Path of the Customer
- Key Organization Contacts and Escalation Path of the Partner
- Key Organization Contacts and Escalation Path of any Carriers and Vendors

Prior to installation and deployment, a detailed escalation document should be provided on a per service basis (Broadband, MPLS, Cellular). The document should include how to raise tickets with the respective service provider as well as the escalation path. Understanding and validation of the process is critical to ensure timely resolution to issues as they arise. Test tickets should be created with each service, carrier, and vendor prior to production deployment to ensure the process works as expected.

Master System of Record and Documentation

The collection of accurate information associated with a site deployment is crucial from an automation and troubleshooting perspective. The automation of collecting and ingesting this data is critical in order to scale from a deployment perspective. Manual manipulation of data sets from various parties is bound to become stale and contain errors. The Zero Touch Provisioning (ZTP) process should take into account the collection of all site details required to build and support a given customer location.

This centralized data collection should be stored in a documentation repository with version control to provide change tracking and easy data recovery. This should be a single, master document that can be easily consumed by automation products (XLS or CSV spreadsheet would suffice).

Below is a bare minimum recommended list of information that should be collected for each site's installation:

Site Details

- Site Address
- Site Contact (Name and Phone number)
- Site Nickname (Store Number, Code, etc.)
- Days & Hours of Operation
- Local Subnet(s) and Purpose of Each Subnet

Circuit Details

- Carrier
- Circuit ID
- Type
- Upload & Download Speed
- QoS Profile
- MPLS ASN
- VLAN Tag (if any)
- IP address
- Subnet Mask
- Gateway IP
- Modem Vendor
- Modem Model

LTE Details

- Mobile Carrier
- Equipment IMEI
- SIM ICCID
- Mobile Number

Hardware Installed on Site

- Hardware Model(s)
- Hardware Serial Numbers

Equipment Location

It is critical to ensure all equipment is installed in the accurate location to avoid costly truck rolls resolving issues that would be avoidable with proper planning. Installation errors are costly to repair from an on-site resource perspective, but also dramatically delay the activation of services which creates reputational impact as well as time to market incurred costs. The following details should be documented to ensure the installed equipment can be mounted and powered properly.

Identification of physical equipment checklist

- Confirm Site Address & Contact
- Obtain site survey reply including:
    - Rack or wall mount specifications
    - Power specifications including redundancy & battery backup requirements, receptacle interface type
    - Picture of desired location that meets space requirements
        - Provide deployment specific installation guidelines

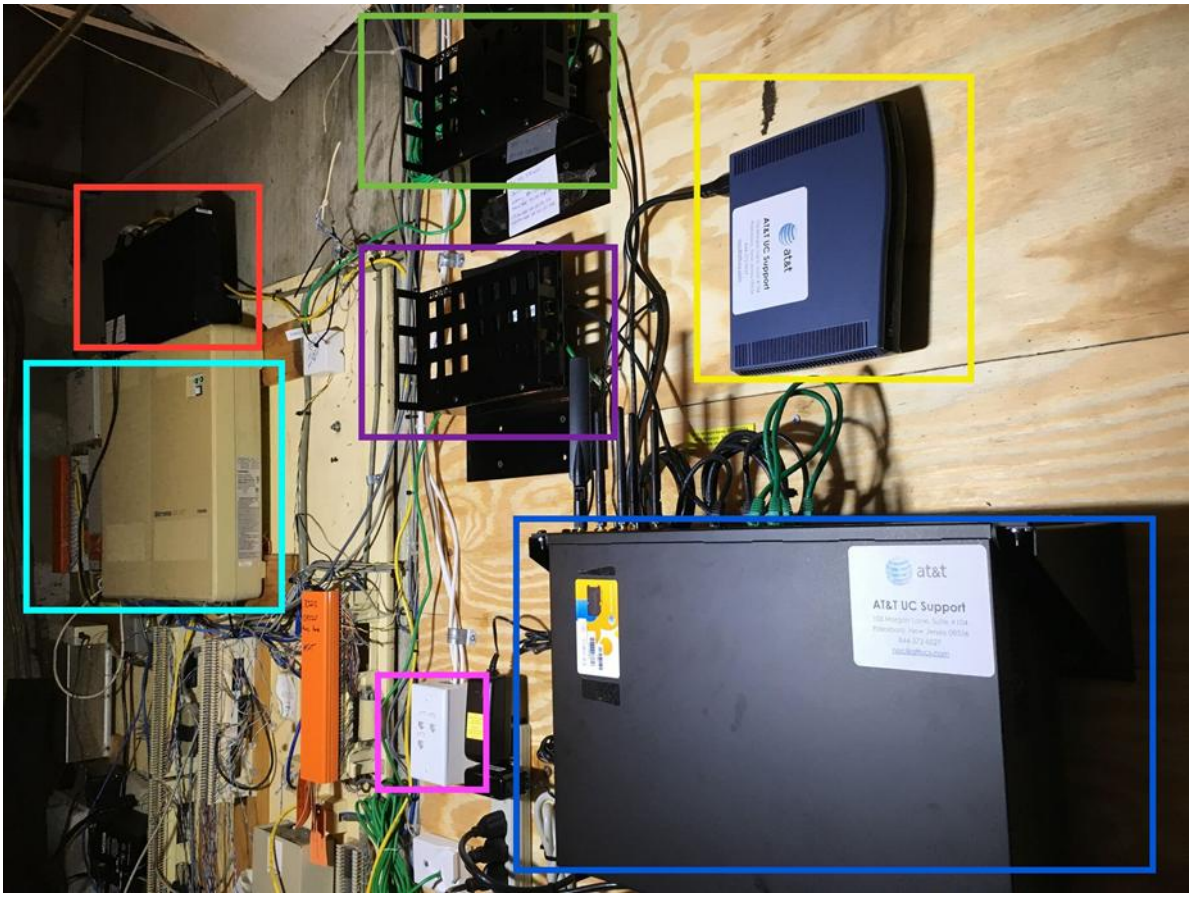

Figure 36: Example Communications Room

In most cases there will be dispatches from several different teams associated with installing a portion of the service (e.g., DSL circuit & modem Install, MPLS circuit and router/switch install, T hardware installation). A picture should be provided to all field service technicians that will visit the site as part of the install illustrating the exact location of the placement of the equipment.

In several instances, circuits and routers have been placed in the wrong location. As an example, a circuit and modem were terminated in a meet-me shared closet servicing an office building. Re-terminating and moving the equipment to the accurate location took over a month to unwind including several dispatches.

The following picture provides the phone board at a given location post-installation. The pre-installation picture would provide the exact location and minimize the risk of installation error. Deployment-specific installation guidelines should be provided as well for the field service technician prior to arrival on site, including rack/shelving specs.

The example provides visibility of the legacy equipment (router and PBX) as well as proper installation best practices.

- Clear labeling of equipment with contact details
- Identification of power receptacles and interface types
- Identification of customer LAN connectivity
- Proper cable dressing and shelving/placement of modems

Notable Exception Process

Not all customer locations will be the same, and some will have notable exceptions. These exceptions *must* be documented, communicated clearly with the customer, and confirmed to be acceptable by the customer. Some examples:

- No UPS Backup or other power backup system present: The customer must be made aware that if the site loses power, the instances will not function during such an outage.
- No power circuit diversity for equipment: If the customer ordered an HA pair of routers, the customer must be made aware that while hardware resiliency is in place, power is still a single point of failure for the solution and the instances will not function during such an outage.
- Other environmental oddities (confined space lacking cooling or ventilation, extension cords utilized instead of power outlets, excessive humidity, etc.) also warrant a notification and confirmation from the customer.

Network Transport Validation

Logical provisioning and ordering details must also be precise to ensure a seamless service delivery process. Re-provisioning circuit, router, and modem logical parameters can be extremely time consuming and add to the delay of the service offering. In addition, inadequate bandwidth from a capacity management perspective will be detrimental to the overall performance of the solution.

- Order validation for all transport services
- Confirm capacity of upload/download speeds
- Accommodate peak existing service throughput in DR scenario
- Accommodate new service throughput in DR scenario (example, Voice over IP, Video, etc.)
- Include 150 Kbps for monitoring & management flows
- Order and confirm logical configuration settings
- Validate system of record for configuration details, required for ZTP
- Obtain vendor escalation documentation on a per service basis (Broadband, MPLS, Cellular)

Capacity Planning

A capacity planning study must be carried out with the customer to ensure adequate bandwidth is provisioned to meet their needs with room for growth. Both upload and download bandwidth

must be considered for this planning. Any additional services being planned as part of the deployment (migration to VoIP, integration of new video services, local hop-off internet connectivity, etc.) must be included with the capacity planning as well.

In addition, Disaster Recovery (DR) scenarios need to be accounted for in order to ensure traffic loads are serviced as per customer expectations. Circuits should be sized to ensure a single circuit can handle the expected load for a given site. Any assumptions made regarding permissions of traffic classes/applications across primary/backup paths should be clearly communicated and documented.

The transport leveraged and associated upload/download speeds must match the expected throughput including monitoring and management flows up to 150Kbps.

Logical Configurations

Tuning of logical settings must be clearly provided and validated upon install as it is time-consuming and cost prohibitive to resolve post install.

Firewall services and TCP/UDP Flood Limit prevention is not required on the modem as that functionality is inherent in the solution which provides a vastly more robust feature set. As such, these features should be disabled.

It is also recommended to ensure public IP Address space is leveraged on the internal and external side of the modem when available. ICMP responses should be enabled on the modem to help triage issues. Finally, Wi-Fi should be disabled on the modem so rogue users are not able to saturate the transport which would be outside of visibility from the perspective of the router.

It should be noted that if MPLS is leveraged as a transport solution, the QoS profile associated with the circuit must match the expected traffic throughput on a per class basis as outlined in the capacity management discovery process or traffic will be dropped.

The following recommendations are provided to ensure modem and carrier settings are set correctly in order to avoid service degradation.

- Confirm logical configuration settings
- Enable Modem public IP address & ICMP response when available
- Disable Modem Firewall
- Disable Modem TCP/UDP Flood Limit parameters
- Disable Modem Wi-Fi
- Ensure MPLS QoS Profiles match expected throughput

Additional Considerations

As part of the deployment process there are a few additional items that should be considered as follow:

Hardware Burn-in

Hardware is most likely to fail within the first 48hrs, or it will last for years assuming it is in a clean room (UPS power, air-conditioned, grounded). All hardware should be tested for a couple of days prior to shipping to its destination at the top end of the operating temperature.

LTE Service

If LTE is leveraged as a transport, the expectations of its use should be clearly identified as described above from a capacity planning and permission perspective (i.e., primary/backup per traffic class, application). Failover and session survivability should be discussed as there are implications to billable bandwidth if sub-second failover is a requirement. The data plan associated with the LTE service should take into account the expected throughput which can significantly increase due to path monitoring when sub-second failover detection is enabled.

LTE signal quality can also be a concern depending upon where the equipment is located. Signal strength should be validated as part of the initial deployment and antenna extension capabilities should be outlined as part of the service offering.

Conclusions & Recommendations

Networks happen. Too many enterprise networks have evolved project by project without an overall architecture. This patchwork of solutions creates technical debt that makes networks complex, fragile, costly, and not agile. Enterprises and government agencies will need to make the initial investment in time and money to adopt SD-WAN that in turn will allow them to become more agile in everything they do.

Network architectures are consistently evolving, because the fundamental principles on how you design networks changes over time. Unlike physics and chemistry that have natural laws that do not change, the underlying technology and business models for networking change every 7 - 10 years. We are undergoing another major industry transition as the underlying switching and routing technology moves to software. Large cloud providers such as Amazon, Facebook, and Google have jumped on this early and have IP networks that are 100x faster at a fraction of the cost of traditional large enterprises or Network Service Providers.

SD-WANs are a major step forward in WAN networking technology. Implementing SD-WAN in and of itself is not a strategy. The most successful SD-WAN implementations are driven by business's transformation of their branch offices, adopting cloud applications, implementing IoT, and supporting mobile users. These implementations are a total transformation in the way networks are procured, built, and operated.

While every SD-WAN will be unique, there are common themes that early adopters are following to ensure success. These include:

1) **Build a New Architecture** – Instead of bolting on SD-WAN to the existing network like they did for WAN Optimization, a successful SD-WAN starts from scratch with no anchors from the existing WAN. This means new network service providers for network access and transport links, new network design that enables users to be within 25ms of the applications they consume today and moving down to 5ms in the future, and new SD-WAN hardware and software. The architecture is vendor agnostic and the SD-WAN software vendor(s) chosen are based on requirements tied to features and costs. The amount of information required to put together a solid architecture is significant.

2) **Use 4/5G Cellular** – Using wireless as a data path for redundancy in case the wireline path is cut or has congestion. Most critical business applications used to run the business are lower bandwidth and transactional in nature. Business applications used in managing the business are higher bandwidth and with voice/video interactions.

3) **Provide End-to-End Networking** – Run SD-WAN from a site all the way to a cloud service provider to ensure end-to-end networking performance and security. Network segmentation plays a big role in taking a physical network and dividing it into separate logical segments. The logical segmentations provide network performance and security capability that can be tied to specific users, devices, services, applications, and data.

4) **Move Towards Zero Trust Security** – This security model does away with a perimeter and trusted internal network. Everything is untrusted, and through authentication and authorization, trust is granted. Whitelist security policies of what is allowed in a least privileged model instead of a blacklist model that defines where network traffic cannot go. Anomaly detection provides early detection if a user is misbehaving or a device has been infected with malware.
5) **Ensure Unified Communications Integration** – With voice being a business-critical application while also being very network sensitive to latency, jitter, and packet loss — and video consuming the bulk of network bandwidth — SD-WAN and UC integration is critical.
6) **Consider a Managed Service Provider** – Implementation speed, experience, and turn-key management is driving most enterprises to seek an SD-WAN managed service provider. MSPs provide an elastic service that can scale up and down in near real-time. In times of crises, elasticity is crucial.
7) **Watch Out for Scale** – SD-WANs with their tunnels and encryption have a lot of overhead, and throughput per box can be limited. Vendors advertise throughput as 10x what reality is in many implementations. All large implementations (thousands of sites and hundreds of thousands of tunnels) have required additional headend infrastructure over the original order.

SD-WANs empower enterprises to transform their networks. First, they are faster in providing greater network performance with direct routing between users and devices to services, applications, and data while being agile and elastic. Second, they are better by providing automation, integration, and being 99.999% reliable to support today's 24/7/365 digital business. Third, they are cheaper by lowering operational expense through use of commodity bandwidth, lower maintenance, and fewer people along with capital expense of using commodity hardware. Lastly, they are more secure by encrypting all data in motion and providing 1:1 micro-segmentation between users and devices communicating with services, applications, and data.

Traditionally, it has taken enterprises 3 - 4 years to upgrade their networks. This is too long in a digital world that requires an enterprise to pivot in near real-time as the market and world changes. Businesses desire "push button" infrastructure that is elastic, which SD-WANs can facilitate once they have been implemented.

About TechVision

World-class research requires world-class consulting analysts and our team is just that. Gaining value from research also means having access to research. All TechVision Research licenses are enterprise licenses; this means everyone that needs access to content can have access to content. We know major technology initiatives involve many different skillsets across an organization and limiting content to a few can compromise the effectiveness of the team and the success of the initiative. Our research leverages our team's in-depth knowledge as well as their real-world consulting experience. We combine great analyst skills with real world client experiences to provide a deep and balanced perspective.

TechVision Consulting builds off our research with specific projects to help organizations better understand, architect, select, build, and deploy infrastructure technologies. Our well-rounded experience and strong analytical skills help us separate the "hype" from the reality. This provides organizations with a deeper understanding of the full scope of vendor capabilities, product life cycles, and a basis for making more informed decisions. We also support vendors in areas such as product and strategy reviews and assessments, requirement analysis, target market assessment, technology trend analysis, go-to-market plan assessment, and gap analysis.

TechVision Updates will provide regular updates on the latest developments with respect to the issues addressed in this report.

About the Authors

Sorell Slaymaker has 30 years of experience designing, building, securing, and operating IP networks and the communication services that run across them. His mission is to help make communication easier, cheaper and more secure since he believes that the more we communicate, the better we are. Prior to joining TechVision Research, Sorell was an Evangelist for 128 Technology which is a routing and security software company. Prior to that, Sorell was a Gartner analyst covering enterprise networking, security, and communications. Sorell also led architecture teams at Target, Travelers, Cigna and United Health.

Sorell is an IT Architect with a focus on network, security, and communications architecture. He specializes in IT Architecture – Network Architecture, SIP Trunking, Contact Centers, Unified Communications, and Security Architecture.

Appendix

VxLAN Overlay Example Layout – Cisco SDA