



Nemertes

SD-WAN as a Service: The Rise of Managed SD-WAN

Maximize WAN Uptime and Minimize Management Overhead with Shift to Managed SD-WAN

John Burke

Principal Research Analyst and CIO
Nemertes Research

Q2 2019



Table of Contents

- Executive Summary3**
- Year of SD-WAN Also Year Managed SD-WAN Arrives4**
- In-Net SD-WAN.....4**
- Drivers and Benefits of In-Net SD-WAN Services.....5**
 - Everyone Has a Solution 5
 - User Satisfaction with Current Offerings 6
 - Last Mile Sprawl: Lots More Internet...Lots More Providers? 6
 - Integrated Cloud Connectivity 7
 - Operational impacts: Higher Cost, Lower Downtime 8
 - Branch Stack in the Cloud..... 8
- Loss of Control.....9**
- Conclusion.....10**
- Methodology.....11**
 - Participants by Title 11
 - Participants by Company Size 11
 - Participants by Technology Culture 12
 - Participants by Industry..... 12

Executive Summary

SD-WAN is rapidly reshaping the WAN landscape and the enterprise's idea of what a WAN should be and can be. More than 40% of organizations participating in Nemertes Research's [WAN Economics and SD-WAN: 2018-19 Research Study](#) have begun to deploy. One factor accelerating adoption is the rapid profusion of managed solutions. While 69.9% deploy a do-it-yourself (DIY) solution, acquiring and operating the SD-WAN gear themselves, already 30.1% use a managed SD-WAN service, most of which deliver some SD-WAN functionality via the providers own service cloud. This represents a huge jump in SD-WAN service adoption in a short time.

Driving this rapid rise in in-net adoption are the spreading availability of solutions from all major service providers, as well increasing preference for as-a-service-based consumption in general. More organizations are also aware of the “provider sprawl” issue—the cost and complexity that come with proliferation of last-mile connectivity providers in an SD-WAN. And, interest is increasing in integrating exchange-based cloud connectivity into the SD-WAN seamlessly. Add to these the operational, agility, scale, flexibility, and security benefits that can come with in-net SD-WAN and a “stack in the cloud” and the momentum of in-net offerings makes sense.

Everyone deploying SD-WAN expects to get improvements in WAN availability and stability: they are foundational promises of the technology. Unsurprisingly, the differences are stark when comparing the overall experience (including a great many organizations just beginning their transitions) with that of those nearing the end of deployment (having 90% or more of their sites converted to SD-WAN). Somewhat more surprising, though, is the dramatic difference between DIY and in-net experiences within the 90%+ group: In-Net services are providing a substantially better experience, resulting better than 90% reductions in overall downtime and in average outage duration.

IT professionals should:

- Evaluate multiple SD-WAN options, both DIY and In-net
 - Do you want just to simplify WAN management while reaping the benefits of SD-WAN, or do you also want to outsource management and get the benefits of “stack in the cloud” technology?
 - Look to your existing WAN connectivity and WAN managed services providers as you evaluate in-net offerings—if they have done well for you so far
- Evaluate your WAN geographic profile and determine how likely you are to have to deal with last-mile sprawl, and to have multiple connectivity options at each site
- Determine your organizational taste for collapsing branch stacks into an SD-WAN solution in a provider cloud
 - Key stakeholders outside the networking group, such as security and risk management, should have input into the decision and selection.

Year of SD-WAN Also Year Managed SD-WAN Arrives

SD-WAN is rapidly reshaping the WAN landscape and the enterprise's idea of what a WAN should be and can be. More than 40% of organizations participating in Nemertes Research's [WAN Economics and SD-WAN: 2018-19 Research Study](#) have begun to deploy the

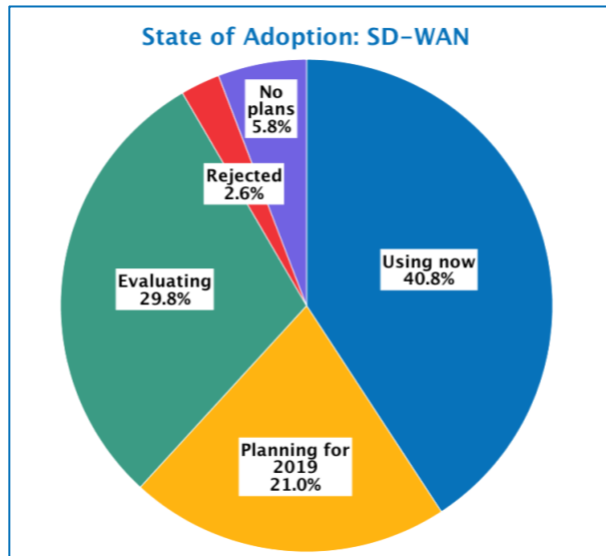


Figure 1: SD-WAN Deployment Begun in 40% of WANs

technology despite its relative novelty, it having only become a well-defined solution set in the last 3 years. (Please see Figure 1.) Most will complete their rollouts by 2021. Indeed, its spread is so rapid, and the infusion of SD-WAN functionality into WAN solutions so broad, that soon the cluster of functions we currently call SD-WAN will be considered table stakes—the default expectation for any branch network.

One factor accelerating adoption is the rapid profusion of managed solutions. While most organizations—69.9%—deploy a do-it-yourself (DIY) solution, acquiring and operating the SD-WAN gear

themselves, already 30.1% use a managed SD-WAN service. This represents a huge jump in in-net SD-WAN adoption in a short time. In [Nemertes Cloud and Networking: 2017-18 Research Study](#), only 8.5% of organizations used an SD-WAN service. Enterprises with larger WANs are more likely to be using a managed offering: 46.2% of those with 100 to 500 locations, and 40.0% of those with more than 500 locations.

Some early adopters (and fewer of more recent ones) procured a managed version of what they could do themselves. Increasingly, though, those deciding not to take the DIY path instead get a service incorporating functionality specific to the provider's own network—an in-net SD-WAN.

In-Net SD-WAN

In contrast to do-it-yourself deployments and managed deployments of DIY solutions, in-net SD-WAN ties SD-WAN functionality to a providers' network service cloud. SD-WAN functions may all be delivered in the service provider's edge and core infrastructure, with a branch using a traditional router to connect to the provider's nearest point of presence; in essence an SD-WAN-equipped Network as a Service (NaaS) solution.

More often, though, some or most functions are provided on-premises via physical or virtual appliances under service provider management. This pushes some work out of the service provider's infrastructure, but more importantly also allows optimizations such as

compression on the last-mile connectivity and greatly eases link load balancing across connections from diverse carriers at the site.

In-net SD-WAN is sometimes tied to Network Functions Virtualization (NFV), with the various functions of the SD-WAN solution provided by separate, cooperating Virtual Network Functions (VNFs) dynamically downloaded to the on-premises device (where there is one) or chained into the traffic path in the carrier infrastructure. These solutions currently rely on the provider to deliver the Customer Premises Equipment (CPE) to run the VNFs, but create the possibility of the CPE being customer owned and generic rather than bespoke for the service.

Drivers and Benefits of In-Net SD-WAN Services

Several market and solution imperatives are driving the rise in in-net adoption. First and foremost, there is the availability of solutions, of course, and the ability of those that prefer as-a-service-based consumption to get an SD-WAN using their preferred model. Beyond that there is also a rising awareness of the “provider sprawl” issue around last-mile connectivity, and a spreading need to integrate cloud connectivity into the WAN seamlessly. Add to these the operational, agility, scale, flexibility, and security benefits that can come with in-net SD-WAN and a “stack in the cloud” and the momentum of in-net offerings makes sense.

Everyone Has a Solution

Every Tier-1 and Tier-2 carrier and cable company has an SD-WAN solution. Most have now been available long enough that the provider has ironed out the major kinks in quoting, planning, and executing on a deployment, has sufficient trained staff, and has an infrastructure ready to support the service. Providers have been keen to bring solutions on line both because of the drumbeat of interest from enterprises, and to capitalize on the very same manageability and service-hardening features that drove that initial wave of intense enterprise interest in SD-WAN. With SD-WAN in place, providers can deliver a better service while also bringing down their own deployment and management costs by reducing errors in configuration and speeding remediation of problems. SD-WAN solutions also let them clearly and easily see per-application performance, link health, and site uptime across diverse connection types, paths, and providers, an unprecedented improvement in network usage and performance visibility.

Many in-net offerings have at their base solutions that, although capable of being deployed DIY-style, were in fact designed to be provider friendly: to be deployed as managed solutions platforms. Among other things, they typically enable through fine-grain APIs the easy development of customer-facing, multi-tenant portals.

Providers use customer-facing portals for visibility, for service requests and trouble reporting, and sometimes for hybrid (shared) management of the solution. The more provider-friendly the underlying solution, the more likely it is to support fine grain separation of access to management features and assignment of permissions, giving the provider the option of defining a broader spectrum of hybrid service options. It might support solutions managed solely by the provider, as well as solutions allowing the

customer to adjust application priority assignments (and nothing else), as well as solutions allowing customers to modify security policies and add new priority groupings, and so on.

User Satisfaction with Current Offerings

On a 0-to-10 scale, users asked to rate their provider based on how likely they were to recommend them to someone else gave the edge to DIY solutions, which averaged 8.4 out of 10 vs. managed solutions' 7.9. Overall, either option is a solid choice in the view of users.

When further broken down into detractors (those rating their provider below a 7) and promoters (those rating them 9 or 10), we see that 49% of DIY users are promoters vs. only 7.8% as detractors, leaving a net promoter score of 41.5%. For managed options, though, only 13.6% were promoters vs. 4.5% detractors, for a net score of 9.1%. As early deployments of older technologies age out of the market and built-for-purpose solutions

come to dominate, we expect overall satisfaction with managed solutions to rise.




Figure 2: Majority Want Others to Manage Last Mile

Last Mile Sprawl: Lots More Internet...Lots More Providers?

SD-WAN's business case often revolves around the hard-dollar economics of connectivity: substituting low-cost Internet bandwidth for high-cost MPLS to meet the burgeoning needs of new communications, collaboration, and cloud applications. In theory, IT can use the lowest cost provider at each location to bring costs down as much as possible, and in a worst-case scenario wind up with more providers than locations.

An explosion in the number of providers has real costs. They include both administrative costs, as the company signs and tracks a contract for, and manages billing for, each one; and operational costs, as IT has to manage the service relationship with each, including the coordination of support processes and staffs.

The problem of provider sprawl can be significant, and is on the minds of SD-WAN users looking for a simpler WAN experience in the future. More than 58% of organizations plan to outsource the problem of coping with the last mile, whether by using last-mile connectivity service aggregators to rein it in (34.7%) or by using a managed SD-WAN provider that will handle that as well as the SD-WAN itself. In-net providers do not all offer management of last-mile links. Some that do make it a part of the service, while others treat it as a line-item

addition to the base service. Some allow a BYO approach and manage any connectivity the enterprise brings with it; others will only support connections with a specific set of ISP partners.

State of Adoption: WAN-Cloud Exchange

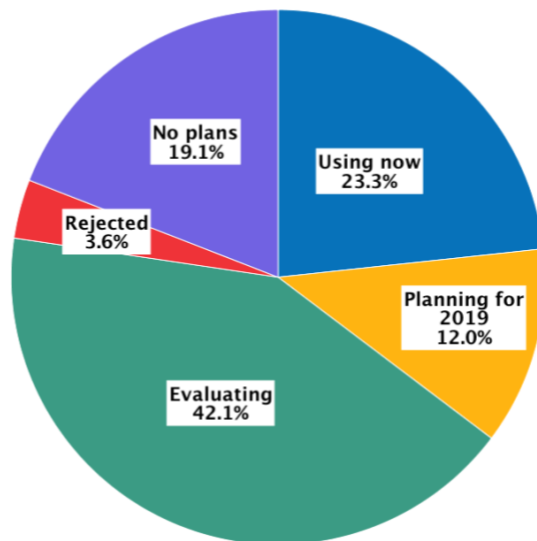


Figure 3: WAN-CX State of Adoption

Integrated Cloud Connectivity

An emerging driver for adoption of in-net SD-WAN is rising interest in cloud exchanges, which allow cloud-bound traffic streams to reach cloud resources more directly. A WAN-cloud exchange (WAN-CX) uses optimal placement of network points-of-presence to allow cloud-bound traffic to exit the provider network in a place where it can enter the Cloud Service Provider (CSP) network directly. That is, the exchange connects to the CSPs, enterprises connect to the exchange.

Using an exchange provides many of the benefits of a direct cloud connect (which links enterprise infrastructure directly to cloud service provider

infrastructure, with a dedicated strand of cable in a shared “meet me” facility). Traffic avoids traversing the Internet and so has more predictable and steadier performance, lower latencies, lower packet loss, and higher throughput. Many organizations also feel their traffic is more secure the less time it spends traversing the public Internet.

To those benefits, using an exchange adds flexibility and agility: flexibility by allowing any size of connection through the exchange to a provider, agility by allowing connections to be spun up, spun down, and resized on the fly via a management portal. So an enterprise, instead of spending days or weeks getting a dedicated, private physical connection into, say, Microsoft Azure at 1Gbps even though only a fraction of that bandwidth is needed, and then again to connect into AWS, and again for Salesforce, and so on, could get 1 Gbps into a WAN-CX and then spin up virtual 10Mbps or 100Mbps links into the clouds of interest. And be able to add a new destination in minutes; or drop or resize an existing one with a few mouse clicks.

Of course, all this comes at a cost, and the cost is...cost. Using an exchange is typically more expensive on a per-Mbps basis than using a direct connect, but for lower bandwidths that is more than made up for by using far less bandwidth. Just shy of a quarter of organizations overall use an exchange currently. (Please see Figure 3.) A significantly higher proportion—43.7%—of those using SD-WAN also use a WAN-CX. All major in-net SD-WAN providers also have WAN-CX solutions, and as use of both spreads, use of the two together will as well.

Operational impacts: Higher Cost, Lower Downtime

Because it is a managed solution, we expect that those using in-net SD-WAN will have fewer WAN staff per site than those going DIY (and they do, .3 vs .41) and will spend more per site for WAN (and they do, \$27916 vs. \$19111). Also as expected with a managed service, DIYers do more work to resolve problems; they spend 18% of WAN staff time troubleshooting WAN issues; In-net users, only 8%.

Everyone deploying SD-WAN expects to get improvements in WAN availability and stability: they are foundational promises of the technology. Unsurprisingly, the differences are stark when comparing the overall experience (including a great many organizations just beginning their transitions) with that of those nearing the end of deployment (having 90% or more of their sites converted to SD-WAN). Somewhat more surprising, though, is the dramatic difference between DIY and in-net experiences within the 90%+ group: In-Net services are providing a substantially better experience.




Figure 4: Differences in Service Resilience, In-Net vs DIY

Branch Stack in the Cloud

In-net SD-WAN allows maximum service delivery with minimum Customer Premises Equipment (CPE) because functionality can be delivered in the service provider cloud, anywhere from edge to core, as well as in any CPE used. Providers can dedicate resources to SD-WAN work as needed, and wherever it is most effective and economical. Some jobs will best be handled on the customer-facing side of carrier points of presence (their network edge), such as packet replication or dropping or traffic compression. Others tasks may be best executed in public clouds or the provider’s core, such as traffic and security analytics and intelligent route management.

People want a lot out of their SD-WAN solution: routing, firewalling, and WAN optimization, for example. (Please see 5.)




Figure 5: Many Roles for SD-WAN (PLACEHOLDER; JEB WILL REDO IN DIFFERENT FORMAT)

Enterprises working with in-net SD-WAN are more free to use resource-intensive functions without feeling the limits of the hardware at each site. They are free to try new functions more frequently and to deploy them more broadly without needing to deploy additional or upgraded hardware. These facts can allow a much more exact fit between services needed and services used, since there is no up-front investment needed to gain the ability to test an added function.

Enterprises are also able to deploy more rapidly. On trying new functions at select sites and deciding to proceed with broader deployment, IT can snap-deploy to the rest. On lighting up a new site, all standard services—as well as any needed uniquely at that site—can come on line immediately, anywhere.

The provider gets more agile too. Working with a software-defined service cloud, it can spin up new service offerings in a fraction of the time required when functions depend on specialized hardware. Rapid evolution of services, as well as the addition of new ones, makes it easier for an enterprise to keep current and to get the benefits of great new ideas.

And, using elastic cloud resources for WAN security functions decreases the load on enterprise networks, and on data center security appliances. Packets that get dropped in the provider cloud for security reasons don't consume any more branch or data center link capacity, or firewall capacity, or threaten enterprise resources. This reduces risk for the enterprise over all.

Loss of Control

The trade-off for handing off management of the SD-WAN is some loss of autonomy as compared to a DIY solution. In the DIY model, the enterprise controls the platform, how it is used, and the connectivity plugged into it. As in any managed service, someone else has their hands on the controls with in-net SD-WAN, though hybrid management models are increasingly available as well. And, depending on the provider and the solution, some in-net solutions limit or eliminate choice of connectivity provider. Since by definition some SD-WAN function is delivered via the provider's network, sites all have to connect to that if only

virtually. Beyond that, some providers allow BYO connectivity from other providers, and some don't.

If you have trouble getting connectivity to all your sites from a single provider, this can become a significant issue, and considerably erode any anticipated savings from shifting connectivity at the branch. Likewise, if you want to have provider diversity for your branches, as well as path and link-type diversity—that is, if you want each branch to have at least two links and from two different providers, e.g. one for MPLS and a different one for Internet. A pure NaaS-style SD-WAN provider, which typically doesn't care how you get traffic to their nearest point of presence since their focus is on the “middle mile” connecting the PoPs, wouldn't get in the way of selecting the edge connectivity you want. A provider with a big conventional last-mile connectivity business might.

Conclusion

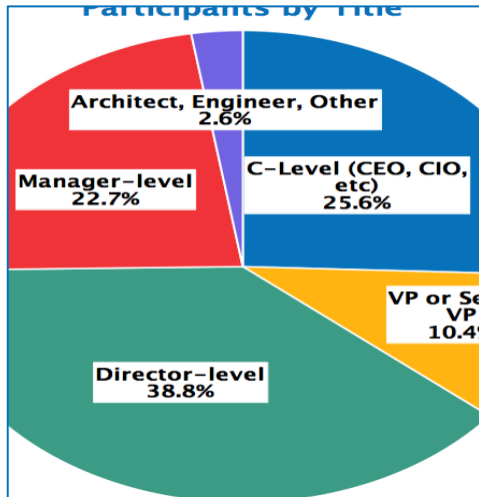
In-net SD-WAN is coming on fast as both products and service offerings continue to mature, and the market for SD-WAN functionality expands. Managed SD-WAN brings the highest degree of ease of use, offering a simple path to eliminating the problem of managing last-mile sprawl, and reducing the amount of time WAN staff troubleshooting WAN issues. In-net services also offer the best improvements in uptime and reliability, bringing about the least annual downtime and shortest average outages. All at a cost, of course.

IT professionals should:

- Evaluate multiple SD-WAN options, both DIY and In-net
 - Do you want just to simplify WAN management while reaping the benefits of SD-WAN (in which case DIY is an option), or do you want to transform your WAN while also outsourcing its management and getting the benefits of “stack in the cloud” technology?
 - Look to your existing WAN connectivity and WAN managed services providers as you evaluate in-net offerings—if you are happy with their performance on your existing services
- Determine your organizational taste for collapsing branch stacks into an SD-WAN solution in a provider cloud
 - Router, firewall, optimizer, even WiFi controller can be rolled into a solution
 - Key stakeholders outside the networking group, such as security and risk management, should have input into the decision and selection
- Evaluate your WAN geographic profile and determine how likely you are to have to deal with last-mile sprawl, and to have multiple connectivity options at each site

Methodology

The Nemertes 2018-19 WAN Economics and SD-WAN research study combined interviewing and surveying to collect information from 309 organizations across North America, Europe, and APAC.




Participants by Title

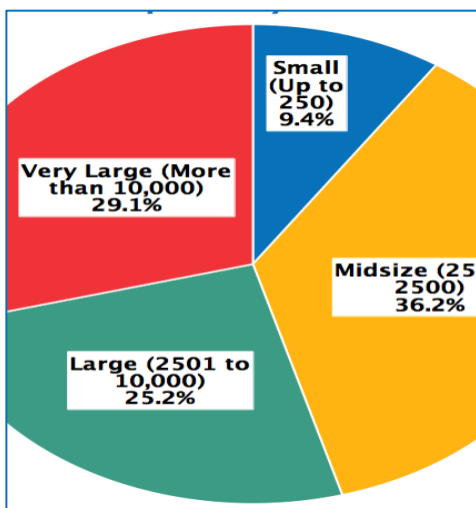
Approximately a quarter of research participants represented senior leadership at their companies. Nearly 50% come from the Director level or are vice presidents.

Participants by Company Size

Nearly half of participants come from companies that are small when measured by annual revenues, bringing in \$300M per year or less. The proportion is significantly higher for some geographies than others. A little more than 28% are large (with revenues greater than \$1B but less than \$10B) or very



large (with revenues of \$10B or more) enterprises.


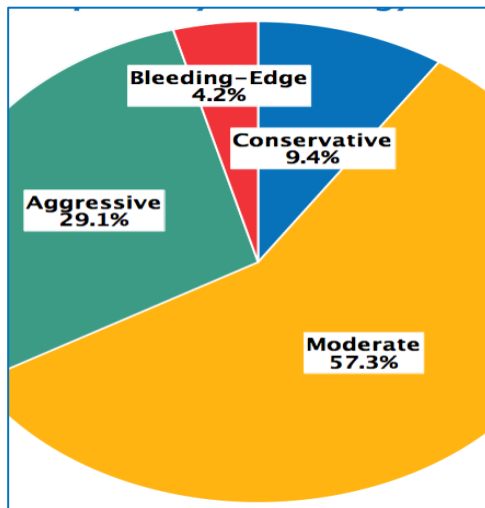


When judged by staff size, a majority, 54.3%, were large (2501 to 10,000 on staff) or very large (more than 10,000 on staff). The bulk of the remainder, 36.2%, were midsize companies having 251 to 2500 on staff.

Participants by Technology Culture

Nemertes defines technology culture as an organization’s overall stance on technology, and asks participants to say which statement best describes their organization’s culture:

- **Conservative:** We deploy technology only when it has been proven to deliver a benefit (usually financial), and generally *after* it has been widely across other organizations. IT is not strategic.
- **Moderate:** We are generally conservative, but make exceptions on a case-by-case basis for specific technologies. IT is somewhat strategic.
- **Aggressive:** We view technology as a competitive advantage and seek to deploy it ahead of most other organizations. IT is strategic.
- **Leading/Bleeding Edge:** We view technology as a competitive advantage and strategic differentiator and deploy it 12-18 months ahead of our competitors to create and sustain our competitiveness. IT is highly strategic.



Participants by Industry

Participants come from more than 17 industries, which can be grouped into the categories shown here. For example, “Financial Services” incorporated several subcategories, such as “Banking” and “Insurance.”

About Nemertes: Nemertes is a global research-based advisory and consulting firm that analyzes the business value of emerging technologies. Since 2002, we have provided strategic recommendations based on data-backed operational and business metrics to help enterprise organizations deliver successful technology transformation to employees and customers. Simply put: Nemertes’ better data helps clients make better decisions.