Turning the WAN Inside Out

COVID-19 has only accelerated a shift already under way. Most WAN traffic touches the outside world, so what does WAN even mean now?

John Burke

Principal Research Analyst and CIO

Nemertes Research

Q3 2020

Table of Contents

Executive Summary

The WAN as we knew it is over. Just 39% of enterprise WAN traffic originates from, and terminates on, enterprise premises. The remaining 61% either originates from an off-premise site (such as a home-office), terminates on an off-premise location (such as an IaaS, PaaS, or SaaS cloud workload) or both (remote office to cloud).

The change in WAN use creates a need for a revised understanding of the WAN. It is no longer the network connecting company sites to each other. It does that, but as a consequence of fulfilling a broader purpose: the WAN is now "the network we control that connects things we care about when they are not co-located."

What does this mean to enterprise technologists? The short version: they need to plan for a wholesale re-architecture of their WANs. This means rethinking everything: WAN technology, including making an SD-WAN riding atop a mixed and changing pool of connectivity the focus of WAN planning; Internet use, both in terms of embracing it as primary connectivity and in terms of Internet-avoidance tools such as direct cloud connects and WAN-cloud exchanges; and security, especially the shift of focus from "edge" firewalls to zero-trust and the software-defined perimeter.

Bottom line: The WAN is undergoing a fundamental paradigm shift. Enterprise networking professionals must change their thinking, planning, and practices accordingly.

IT professionals should:

- Review WAN strategies and roadmaps to root out plans and policies built on the false assumption that the WAN is for connecting inside users to inside resources
- Evaluate an SD-WAN solution (DIY or managed) to replace their existing WAN, or evaluate whether to continue with the solution they have already, with an eye towards its ability to serve as the control nexus for all external communications as well as for inter-site traffic.
- Engage security teams around the concepts of software-defined perimeter and zero-trust, and the role the WAN can play in making them real.
- Explore the available complementary solutions for connecting to and among clouds, including DCCs, WAN-CXes, and cloud networks.

"I Do Not Think That Word Means What You Think It Means"

When the WAN was invented, it had a single and clear purpose: connecting the different sites of an organization to each other. It accelerated the infusion of technology into back-office operations, then front offices, and ultimately every aspect of operations. It drove the concentration of IT work in data centers by enabling client-server applications to serve any user in any location. And as IP WANs matured, they gradually absorbed the functions of other, specialized networks, especially for voice and video communications.

However, since the widespread embrace of the Internet for business in the 2000s, the corporate WAN has had to accommodate traffic ultimately bound for or starting from places outside the enterprise itself. IT teams have adapted their thinking, but at the core, the WAN is still the "inside" network, and other use cases are exceptions—side gigs. This service to outside destinations was essentially an add-on to the prime use case, since all the real work of the enterprise got done in the data centers or branch office server rooms. One swapped email with folks in other organizations, and researched things on others' web sites, but the real action was still within the company's borders.

However, that balance shifted again, driven by the creation and adoption of software as a service and infrastructure as a service. With SaaS and IaaS, key enterprise IT workloads run outside the enterprise's walls, on infrastructure the enterprise neither owns nor controls. When Nemertes began collecting data on cloud use in 2009 only 2% of organizations were using IaaS and 42% using SaaS. Only a few percent of the overall IT portfolio was running on external clouds. A decade on, in 2019 the average company was for the first time running the majority of its IT work in external clouds; only 40% remains in the data centers.

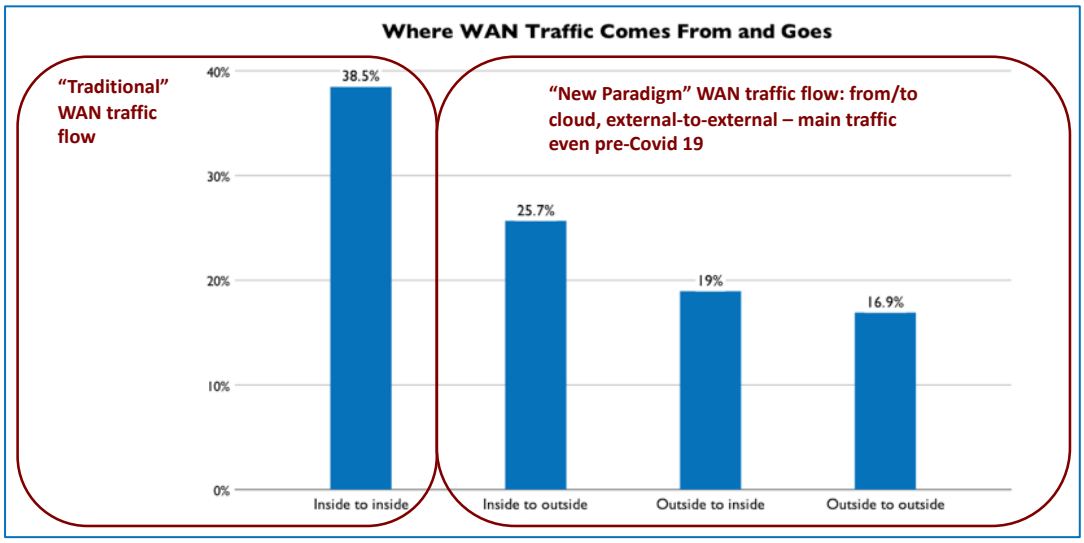

It's not so surprising then, to find that only 39% of the average WAN's traffic was running inside to inside immediately pre-COVID-19. The subsequent massive shift to work-from-home for whoever can has accelerated this change, and the fall back post-COVID will not be to the *status quo ante* but to a new, higher level, both because some initially temporary WFH arrangement will be made permanent, and because the shift to WFH is accelerating shifts of workloads to the cloud, as well.

With the majority of the WAN's job now and in the future being to accommodate traffic starting or ending outside the enterprise's physical boundaries, it's time to ask: what do we mean when we say "WAN?" More importantly, what should we mean?

I Love the Smell of a Paradigm Shift in the Morning: From WAN as Places to WAN as Power

The fundamental change in WAN thinking, the foundational shift in the WAN paradigm, has to be giving up any connection to the physical incarnation of the enterprise. The WAN may serve an enterprise's sites (if it has any) but that does not define what it is or all it does.

Instead, it is crucially important to think of the WAN as "the network we control that connects things we care about that are not co-located." It may be connecting stores to each other and to the cloud, or it may be connecting home offices and mobile workers to data centers, or it may be connecting virtual data centers within clouds to each other. Nearly infinite variations are possible! The thing that defines and limits the scope of the WAN is control: the enterprise has control over what traffic goes where. This includes controlling ingress of traffic coming from sources outside enterprise control, and egress to same.

As enterprises have gradually eroded their borders by moving work out of data centers and into clouds, and as they have interpenetrated their own systems with those of partners, customers, suppliers, and service partners, old WAN thinking struggled to stretch to accommodate. The WAN-as-locus-of-control model makes it as simple to face the disappearance of the border as it is to embrace the waning importance of the traditional site-to-site function.

Magic Carpet Ride: Technologies to Support the Whole New World of WAN

Clearly, a WAN uncoupled from physical sites still needs some kind of infrastructure. However, in the next generation WAN this means mainly virtualized infrastructure running on generic silicon, either network chips or standard x86 compute nodes, rather than custom silicon with tightly coupled operating systems. That is, the new WAN looks more like the data center or a cloud environment.

SD-WAN

Software-defined WAN, bringing the primary tenets of SDN to the WAN, provides the core of the new enterprise WAN. Using centralized, policy-driven management, it allows application-, location-, and sometimes user-sensitive control over how traffic is allowed to flow among endpoints (physical or virtual, wherever located) and how it is allowed to enter from or exit to the Internet. We see nearly two-thirds of companies (61.7%) having begun their deployments, most of which will take a year or more to complete. Another 14% plan to begin deploying before the end of 2021.

One key change we have seen in enterprise use of SD-WAN in the last few years has been an expansion beyond legacy use cases (connecting branches to each other, or the data center, or the Internet) to back-end-focused interconnections: data center to data center, data center to cloud, and cloud to cloud. A majority now use SD-WAN in those scenarios, or plan to within the next 18 months

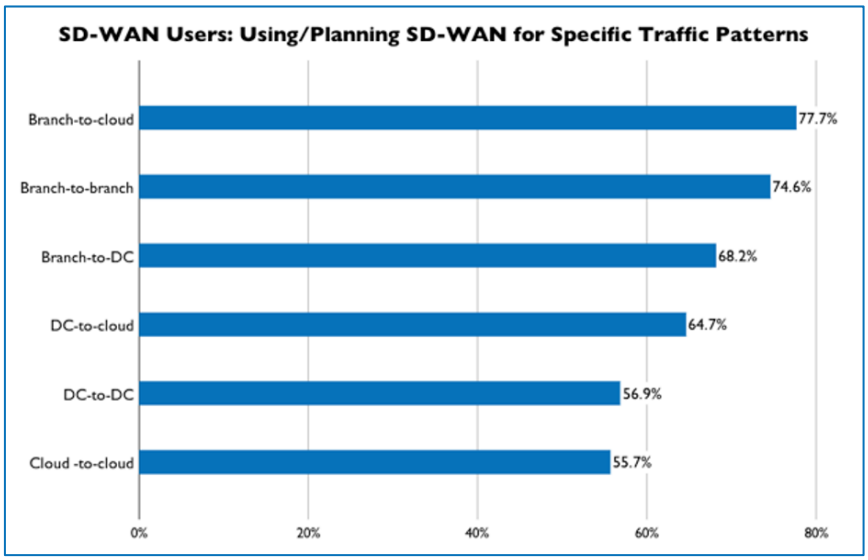

This brings all those pathways under the same policy control mechanism, simplifying the challenge of getting policy consistent across them all.

Direct Connection and WAN-Cloud Exchanges

One function of an SD-WAN is to optimize performance for applications, and a big and growing piece of that is to manage the egress and ingress of Internet and especially cloud-bound traffic. Two technologies enterprises deploy more frequently in the next-generation WAN bypass the Internet for cloud access: direct cloud connects (DCC), and WAN-Cloud exchanges (WAN-CX).

Direct Cloud Connects: My edge connects to cloud edge in meet-me facility

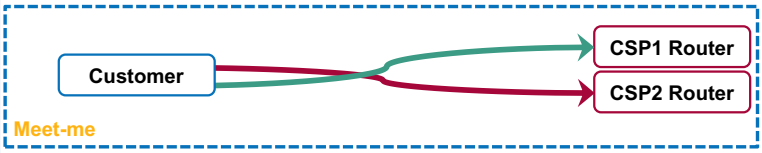

- 33.4% use
- 10.6% plan to by end of 2021

- 31.4% use
- 11.2% plan to by end of 2021

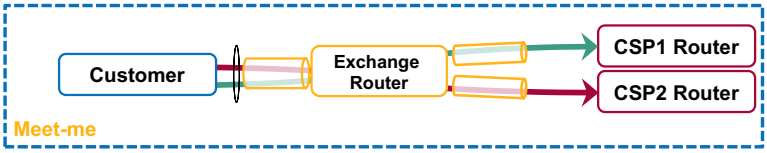

WAN-Cloud Exchanges: My edge connects to exchange in meet-me facility, exchange connects to many clouds' edges

The former connects an enterprise's WAN directly to the cloud service provider's network; the latter connects the WAN to an exchange service, which is in turn connected directly to

multiple cloud service providers, allowing enterprises to spin up virtual DCCs to any of those providers.

DCCs and WAN-CXes expand the SD-WAN's range of options for application traffic paths by adding ultra-low-latency and -loss options. The SD-WAN can then evaluate which path is going to give the best performance for a given packet stream and choose accordingly. It may route all traffic generated in a data center and heading to AWS out through the DCC or WAN-CX serving AWS, for example, but send a remote branch office's AWS-bound traffic through the Internet instead, if Amazon's on-ramp is closer and performing well enough.

Cloud Networking

As noted above, SD-WANs are increasing being deployed to serve as conduits for cloud-to-cloud communications. Cloud-delivered networking solutions provide an alternate, complementary approach to that problem, analogous to using DCCs and WAN-CXes to supplement Internet access to clouds. SD-WANs can connect to cloud-delivered networks, handing off relevant inter-cloud traffic, or possibly use the cloud network to help deploy and manage the SD-WAN endpoints.

Software Defined Perimeter

Conceptually, a software-defined perimeter (SDP) is a detailed map controlling what kinds of network traffic flows are allowed, based on the identity of the entity communicating (whether human, software, or hardware). It is a zero-trust architecture, meaning that if traffic is not specifically sanctioned by that trust map, it is not allowed to flow: a network entity with an SDP around it will only be sent traffic it is supposed to see. Enterprise security is moving towards zero-trust broadly: already, 34% of organizations have begun deploying a zero-trust security architecture, and another 35% plan to by the end of 2020, though it will be a multi-year effort in almost all cases.

An SDP solution consists of a database describing that trust map and a set of policy enforcement points used to control the flow of traffic in accordance with it. SD-WANs and SDPs are, potentially, natural partners, with the SD-WAN able to serve as a distributed enforcement point helping instantiate the SDP. Ultimately, though, SD-WANs will have to evolve towards themselves being an SDP for communications among and outside enterprise-controlled locations, in order to mesh seamlessly with those controlling campus LAN and data center networks.

Conclusion

Even before COVID-19, the legacy WAN was no more; COVID-driven WFH has just pushed that much further behind reality. IT leadership needs to let go of the mental model of the WAN they carry around, based on that legacy WAN, end embrace a new paradigm: the WAN is not a private network linking a bunch of company buildings, it is a locus of control for the connections among enterprise users and resources.

With that in mind, SD-WAN is the logical organizing concept and technology, and along with complementary solutions like direct cloud connects and cloud exchanges, cloud networking and software-defined perimeter solutions, it will form the core of the next generation WAN.

IT professionals should:

- Review WAN strategies and roadmaps to root out plans and policies built on the false assumption that the WAN is for connecting inside users to inside resources.
- Evaluate an SD-WAN solution (DIY or managed) to replace their existing WAN, or evaluate whether to continue with the solution they have already, with an eye towards its ability to serve as the control nexus for all external communications as well as for inter-site traffic.
- Engage security teams around the concepts of software-defined perimeter and zero-trust, and the role the WAN can play in making them real.
- Explore the available complementary solutions for connecting to and among clouds, including DCCs, WAN-CXes, and cloud networks.