Dell'Oro Group

Network Considerations in the Age of Secure Access Service Edge (SASE)

Enterprise networks are approaching a major tipping point, driven by the shift from employees working at a corporate office to working from anywhere. Enterprises are quickly realizing that legacy network and security architectures are inadequate. They must evolve toward a unified networking and security service that increases scalability, agility, and security in a user and application environment that has become highly distributed and mobile across the Internet. A new era of work from anywhere is upon us.

WHITE PAPER

Document Date: May 2021
Author: Mauricio Sanchez

INTRODUCTION

In our paper, we discuss the key drivers of the work from anywhere trend and why legacy wide area network (WAN) and virtual private network (VPN) approaches are no longer effective. We use customer stories to illustrate the key requirements met by leading enterprises as they shifted toward cloud-based networking and security. We also examine the role played by cloud-based technologies, such as Netskope. The industry has embraced the secure access service edge (SASE) as the name for the architecture and technologies that deliver cloud-based networking and security. We provide our definition of SASE, describe Netskope's participation, and consider why Netskope chose to build its own cloud network, named NewEdge.

Contents

THE NEW "WORK FROM ANYWHERE" ERA

While remote work is not new, the COVID-19 pandemic has fueled the significant acceleration of two trends that underly a new style of remote work that we call work from anywhere.

1

Shift from an on-premises to an Internet-based application infrastructure

For more than a decade, enterprises have been migrating from on-premises to public cloud-based applications and computing models for their businesses. The redistribution of the workforce and the higher degree of online business driven by the COVID-19 pandemic has accelerated the adoption of many technologies that enable enterprises to digitalize their operations and process.

2

Growth of a highly distributed remote workforce

According to Dell'Oro analysis of U.S. government and public research data, at its 2020 apex, the pandemic drove a 450% increase in the number of U.S. employees working remotely full-time or occasionally compared to the pre-pandemic baseline. While rates have started to decline, we anticipate long-term remote work rates to settle at 200% above the pre-pandemic baseline. For the financial, information, and professional services industries, the percentage of employees with the ability to work remotely full-time or occasionally is predicted to remain above 70%.

TRADITIONAL APPROACHES FAIL AS THE NETWORK PERIMETER DISSOLVES

For decades, enterprises favored the hub-and-spoke network topology (Figure 1). The hub—the corporate headquarters and data center—was at the center of the network. The spokes emanating from the central hub were the network's connections to individual branch offices and remote users.

Figure 1: Legacy Network Topology

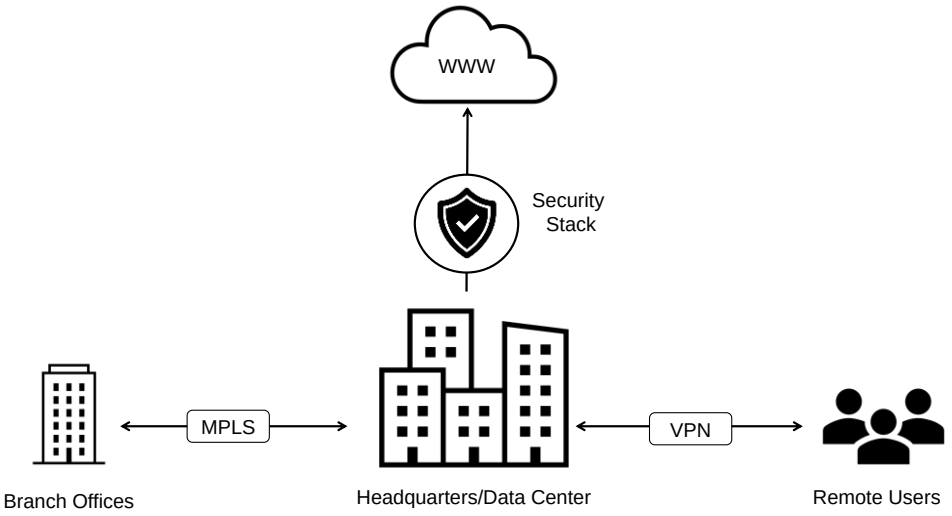

Historically, enterprises have used WANs to connect geographically dispersed locations to each other, to data centers, and, more recently, to cloud-based services. The preferred network technology for WANs has been multiprotocol layer switching (MPLS) due to its reliability and scalability. However, MPLS commands premium prices and is generally difficult to upgrade or expand.

In the legacy network topology, remote users relied on VPN agents. These were based on internet protocol security (IPsec) or a Secure Sockets Layer (SSL) located on each endpoint to connect over the Internet back to VPN concentrator devices at corporate headquarters.

LEGACY WAN AND VPN ARE OUTMODED IN THE WORK FROM ANYWHERE ERA

As the popularity of Software as a Service (SaaS) applications and highly distributed workforces has grown, so, too, have the limitations of the hub-and-spoke topology. There is a new confluence of cost, application experience, and security problems:

High Cost of MPLS

MPLS spending for branch offices became untenable as network bandwidth requirements increased appreciably in order to service SaaS applications that required significantly more bandwidth.

Poor Application Experience

Application experience suffered because traffic from branch offices or remote users to an Internet-based SaaS application first had to be backhauled to the Internet gateway in the corporate headquarters.

Perimester Security Circumvention

As Internet-based SaaS apps became popular, some remote users began to skip the enterprise network altogether by using the Internet to go directly to those apps. This created enormous security blind spots. Thus, it became impossible for IT teams to enforce corporate security.

IT Staff Pressure

Many IT teams find themselves lacking the necessary skillset to manage the new cloud- and mobile-centric IT environment, which further amplifies the cost, application, and security problems noted above.

Over the last ten years, new networking technologies, such as software-defined WAN (SD-WAN), and new security technologies, such as cloud-based secure web gateways (SWGs), were developed. They individually addressed the deficiencies of legacy WANs and VPNs. However, they lacked substantive integration.

In the following section, we describe the reasons why networking and security need to unify.

WHAT ENTERPRISES NEED AND CONSIDERATIONS FOR NETWORK TEAMS

To better understand the needs of enterprises, we interviewed three organizations, two end-users, and a consulting firm, at the forefront of transforming their or their clients' organizations:

- Stuart Walters, Chief Information Officer at BDO UK (United Kingdom), the UK member firm with 6,000 employees. BDO UK is part of the $10 billion international network of public accounting, tax, consulting, and business firms operating under the BDO name.
- Mark Mahovlich, Vice President of Strategy and Execution at ICM Cyber, a cybersecurity consulting firm headquartered in Jackson, Mississippi with 500 clients across 40 U.S. states.
- A senior cybersecurity executive at a global accounting firm with nearly 400,000 employees, who has requested to remain anonymous. This leader is responsible for specifying his company's data-protection strategy across more than 35 global regions.

In the following section, we discuss highlights from our interviews with these leaders.

The unanimous consensus of the group was that legacy architectures are outmoded, inadequate, and overly restrictive. All agreed on the need for a cloud- and mobile-first IT strategy in which security and network are viewed not as individual silos but as a larger, intertwined system. Security and network must be collectively addressed. For many organizations, including some in the group we interviewed, undertaking such significant change is daunting. A best practice voiced by our interviewees was the importance of ensuring strong relationships between, and a sense of collective ownership by the security and network teams.

Mark Mahovlich further explained that while, in principle, networks want to embrace IT transformation, they are still too often hesitant to embrace actual operational change. There is a good reason for their reluctance.

Unlike security teams whose effectiveness is measured by how well they protect data—for example, by blocking traffic to prevent the unauthorized exchange of data—network teams are evaluated based on their ability to transport data from source to destination reliably, on time, and efficiently. The best practices that help ensure that network teams hit their metrics are grounded in decades of trial and error that surround networking as an IT practice. They are biased toward directly

owning—or at least directly controlling—the boxes, wires, and fibers that make up the enterprise network.

As enterprise network architectures become increasingly cloud-centric, Mark has seen how network teams increasingly worry. If network teams no longer directly own nor control the network, how will they perform their job and keep the network up?

Therefore, as the network and security merge in the service of the modern enterprise, the changes that result must be acceptable to the network teams. But beyond mere acceptance, we heard three reoccurring themes from the group about the opportunities to improve networking as IT practice that are opened up by the unification of networking and security:

- The opportunity to **simplify** the network by streamlining network design and deployment, and boosting availability and resilience.
- The opportunity to **modernize** the network by instantiating a WAN edge that secures traffic without performance tradeoffs.
- The opportunity to **optimize** the network by reducing administrative and network costs, while enhancing application performance.

Next, we describe ongoing industry efforts to develop a market for unified networking and security solutions and considerations for the network architect.

ASCENSION OF SECURE ACCESS SERVICE EDGE (SASE)

A recent shift has taken place in how to think about connectivity technologies, such as SD-WAN, and security technologies, such as SWGs. Instead of regarding them as separate tools to be selected and managed, in the current view, the problem and solution space is viewed as a continuum. Now the focus is on developing an integrated platform of connectivity and security services.

What is SASE?

SASE (pronounced "sassy") is a service-centric, cloud-based solution providing network connectivity and enforcing security between users, devices, and applications. SASE utilizes centrally-controlled, Internet-based networks with built-in advanced networking and security-processing capabilities (Figure 2).

Figure 2: SASE Architecture

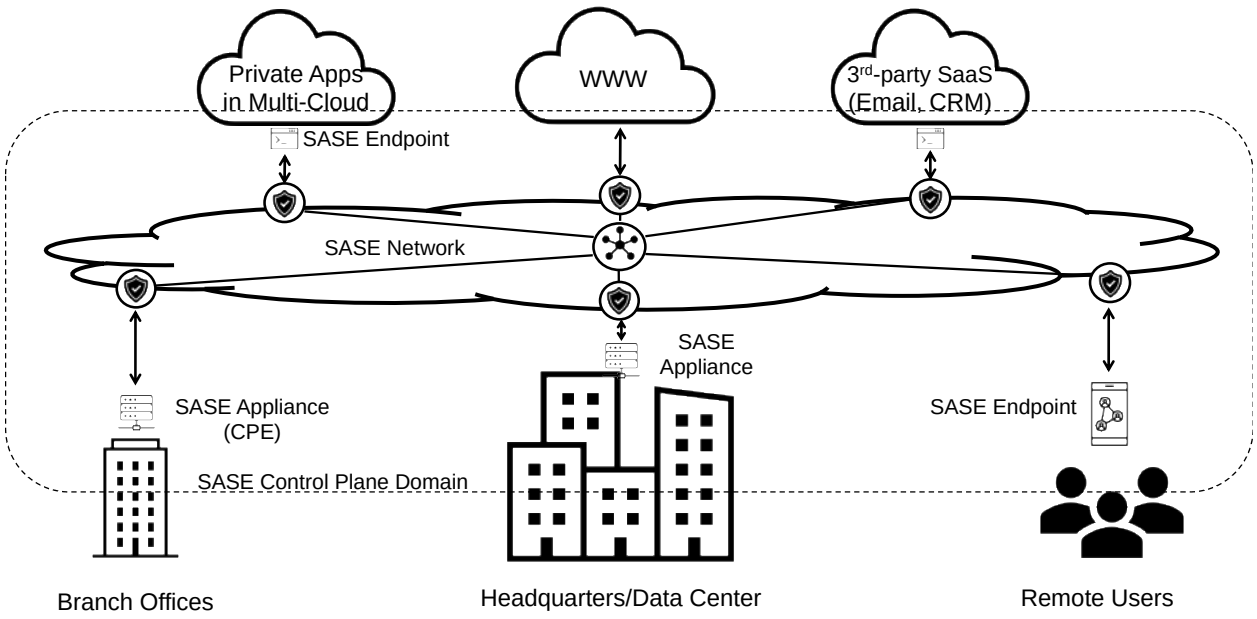

On a tactical level, we view SASE as the union of networking and security technologies that have existed for many years and proven themselves effective in addressing key networking or security pain points. For its networking functionality, SASE heavily leverages SD-WAN technologies. For its security functionality, SASE leans greatly on cloud-based SWG and network-based cloud access security brokers (CASB).

SASE shares the same architectural underpinnings as SD-WAN and cloud-based SWGs and consists of five components:

1

SASE Control Plane Software

The SASE control plane software is the brain of a SASE solution. It makes network and security decisions based on multiple criteria, including user identity, device state, time of day, user/device location, bandwidth, latency, destination, and application.

2

SASE Network

If the SASE control plane software is the brain, then the SASE network is the backbone. Its role is to connect SASE endpoints and appliances and bridge their connectivity to the Internet.

3

SASE Functions

SASE leverages network function virtualization (NFV) to run virtual network functions (VNFs) in a distributed fashion, nominally in the cloud within the SASE network. VNFs perform either a network function or a security function.

4

SASE Endpoints

SASE endpoints include all the users, devices, and applications capable of directly interfacing with the secure network connectivity services provided by the SASE solution.

5

SASE Appliances

SASE appliances optionally process traffic from users, devices, and applications that are unable to connect directly to a SASE network. For example, a SASE appliance may be deployed at branch offices to onboard traffic from all users, devices, and applications to the SASE network.

Netskope as a SASE vendor

We see vendors entering the SASE market based on a core competency in networking, such as SD-WAN, or security – notably SWG and CASB. In our October 2020 analysis, we identified 27 vendors marketing SASE solutions, including Netskope.

Netskope is an example of a SASE vendor entering from a core competency in security, primarily CASB, data leakage protection (DLP), and SWG, and partnering with an SD-WAN vendor, such as Silver Peak (recently acquired by Aruba, an Hewlett-Packard Enterprise company), to offer customers a complete, multi-vendor SASE solution. SASE components that Netskope provides, versus those via partnerships, are shown in the following table.

Figure 3: SASE Components and Netskope

Component	Netskope or Partner	Description
SASE control plane software	Netskope	Netskope Platform
SASE network	Netskope	Custom cloud network: "NewEdge"
SASE functions	Security: Netskope Networking: SD-WAN partners	Netskope provides SWG, DLP, CASB, cloud security posture management (CSPM), and zero-trust network architecture (ZTNA) services
SASE endpoints	Netskope	Windows, Mac, Android, and iOS Netskope client agents
SASE appliances	SD-WAN partners	Integrations with major SD-WAN vendors, such as Silver Peak, VMware, Versa, Cisco, Fortinet, and Aryaka

If there is any SASE component that network teams will scrutinize, it will be the SASE network. In the following sections, we analyze its role more deeply and explore implementation options and key considerations from the perspective of the enterprise network architect. We then examine the customer perspective on Netskope's custom SASE network, which is marketed under the NewEdge name.

THE SASE NETWORK: THE BACKBONE OF SASE

Underlying the SASE network is a function-optimized and security-focused network with most, if not all, of the compute capacity to enforce network and security policy. The SASE network serves as the foundational connectivity and processing substrate in a SASE solution. It must meet three exacting requirements:

#1

Connect enterprise users, devices, and applications anywhere in the world

#2

Undertake network and security processing of traffic

#3

Be highly available and performant without sacrificing requirements #1 or #2

Implementing a SASE network that can satisfy the first requirement is trivial because the Internet already exists and is proven to suitably interconnect the world. Creating a SASE network that addresses the second requirement is becoming trivial, as the global footprint of compute pools available on the Internet grows by the day, whether, in the context of public cloud service providers, content delivery networks, or edge computing providers. However, crafting a SASE network that satisfies all three requirements is extremely challenging for two reasons: the internal structure of the Internet and the finite speed of light.

The Internet is a network of networks with the larger networks run as businesses by communications services providers. Individual networks seek to do good for their users or customers, while being mostly indifferent, and occasionally hostile, to other networks. The realities of the Internet are a significant challenge in setting availability and performance guarantees of any kind.

The speed of light is approximately 186,000 miles per second (about 299,000 kilometers per second). While on a human scale this is blindingly fast, on a machine scale it is not fast enough. Let us consider an example of a U.S.-based SASE solution with its SASE security functions hosted in Los Angeles. Let us then assume a user pops up in Italy seeking to access an application in Germany. With the SASE functions hosted in Los Angeles, this means traffic has to travel extreme distances—from Italy to Los Angeles and then onward to Germany. Application experience will likely suffer when the network distances are so great. Light is just not fast enough.

The path forward to overcome the Internet and speed of light challenges and satisfy all three requirements noted above is to build the SASE network as a highly distributed set of efficient processing capabilities in many locations and networks while minimizing the network distance to users and the applications they access. In this manner of design, a SASE network can both support

the necessary throughput, while also bounding traffic latency. All too often, this is the source of significant grief to network architects and is difficult to mitigate. We see a great focus on SASE networks already taking place, as SASE vendors differentiate in their approach to and in the attributes of their networks.

We see SASE vendors using one or more of the following approaches to implement their SASE network:

- On top of public cloud service providers, such as Amazon Web Services, Microsoft Azure, or Google Cloud
- On top of public content delivery networks (CDN) and edge computing providers
- Building out a custom cloud network

Each approach has its benefits and drawbacks. In the following sections, we describe each in greater detail.

Public Cloud Service Providers (CSPs) as the SASE Network

Global, public CSPs, such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform are obvious choices to host a SASE network. The scale of a public CSP's compute, storage, and networking is unmatched. The approach is popular because of the:

- ***Large number of Points of Presence:*** Public CSPs offer a large number of network entry/exit points spread throughout the world that simplifies getting traffic into a CSP network.
- ***Strong likelihood of application locality:*** In some instances, traffic may not even need to exit if the destination application or service is hosted in the same CSP.
- ***Economies of scale:*** Public CSPs offer opportunities to leverage their massive, elastic compute pools to process traffic and enforce network and security policies. This carries upfront costs and has time-to-market implications. A SASE vendor need not purchase, deploy, or manage the underlying hardware infrastructure. Spinning up new cloud networks is trivial and allows quick entry into the market.

However, there are various pitfalls when implementing SASE networks in public CSPs:

- ***CSPs are oriented to host applications and data, not the network:*** The design intention of the majority of public CSPs is to host customer applications and their data. As such, the emphasis is to be a traffic destination, not an intermediate pathway. This runs counter to a SASE network's end goal: to be a transitory inspector, not the final destination. Moreover, financial hurdles may arise due to the data egress many public CSPs charge when data exits the network.

AWS is the largest public CSP by revenue, according to Dell'Oro analysis. It exemplifies the historical CSP orientation toward hosting applications and data. As of March 2021, AWS has approximately 230 Points of Presence and 25 global regions.[1] Points of Presence allow traffic to enter or leave AWS' network, while regions are collections of one more data center in geographic proximity. For example, in the U.S. there 21 Points of Presence – usually located in major metropolitan areas – and four public use compute regions: Oregon, Northern California, Ohio, and Northern Virginia.

Points of Presence have limited compute capability and traffic must be backhauled to a region's data center for processing. This means a SASE network built on AWS must often backhaul traffic, which may add undesirable network latency and adversely affect application performance — in particular if the Point of Presence is far from the region's data center.

AWS recognizes this limitation of its current architecture and has recently introduced the concept of Local Zones that place compute, storage, and database services closer to end-users. However, by the end of 2021, AWS expects to have only 15 available Local Zones that serve the U.S. only.[2]

- ***Black box challenge:*** Public CSPs have invested heavily in network capacity tools that facilitate the cloud's abstraction of elasticity. However, these tools are not available to customers or they are available only in simplified form. Customers have little visibility or control over the routing, network peering, and traffic tuning, which complicates troubleshooting. SASE vendors, like other CSP customers, may find it difficult to make informed network-planning decisions or troubleshoot problems.

- ***Inability to own destiny:*** The historical focus of public CSPs has been to maximize the profitability of their total customer workload. This, in turn, leads to design and implementation choices that generally favor the CSP rather than individual customers. While some CSPs are beginning to provide optimized variants of their clouds, customers are still at the mercy of their chosen CSP's direction and choices.

Ultimately, implementing a SASE network in a public CSP is possible. However, SASE vendors must invest substantial effort to prevent unacceptable limitations in their solutions.

[1] https://aws.amazon.com/about-aws/global-infrastructure/
[2] https://aws.amazon.com/blogs/aws/in-the-works-more-aws-local-zones/

CDNs and Edge Compute Providers as the SASE Network

Optimizing Internet-based application performance began over twenty years ago with the first wave of CDN services and, more recently, edge computing providers. A CDN implements a performance-focused network and uses the Internet as the underlying transport. Unlike the best-effort interactions among Tier 1 and Tier 2 network providers that compose the general-purpose Internet, a CDN inserts its own management and coordination mechanisms to use the general-purpose Internet selectively. This achieves better results. Edge compute providers extend the concept of CDNs by instantiating distributed compute pools for use by customer applications.

The strong focus on Internet performance bodes well for a SASE network. The end goal of CDNs and edge computing to be the pathway for traffic and not the destination is exactly aligned with the intentions of a SASE network. However, some of the same pitfalls associated with public CSPs also apply:

- ***Black Box Challenge*:** Similar to public CSPs but perhaps to a lesser degree, the entrée of toolsets available to SASE network implementers is likely to be less than what CDN or edge compute providers have at their disposal.
- ***Inability to Own Destiny:*** Similar to public CSPs, a SASE network is at the mercy of the strategic direction and choices made by the CDN or edge compute provider.

While CDN and edge computing providers are in a better position to address SASE network requirements, they still require effort to adopt.

Custom Cloud Network as the SASE Network

Just as CIOs make buy-versus-build decisions for the technology in their enterprises, SASE vendors also have the same choice to make. We consider the two options described above to be analogous to buying technology for a SASE network. The alternative, and the basis for a third option, is to build a custom cloud network. Starting from scratch is not for the faint of heart and comes with its own sets of advantages and drawbacks.

The advantages offered by the custom cloud approach include:

- ***Ability to meet and exceed the three, key SASE network requirements:*** Assuming the SASE vendor has the necessary technical skillset, a custom SASE network cloud should be able to meet, if not exceed, the three requirements we previously outlined.
- ***Controlling its own destiny:*** Being able to make all-important design and implementation decisions, both now and in the future, ensures that a SASE vendor owns their destiny.

- ***Overcoming the black box challenge:*** Though a SASE vendor may not own and control every last element in their cloud network, they do own enough to avoid falling into the black box trap. The following are important control points that the custom approach affords:
  - **Coverage:** Where to instantiate new data centers that serve as points of presence
  - **Breadth:** Ability to ensure that every service is available at every data center
  - **Scaling:** When to add data center capacity to support a scale-out expansion
  - **Capacity:** Allocation of network bandwidth per data center
  - **Performance:** Where to instrument and tune the network to ensure end-to-end visibility and control of latency and roundtrip-times
  - **Resilience:** Ability to automate failover and minimize the blast radius during downtime
  - **Peering:** Choice of with which Tier 1 and Tier 2 networks and public CSP or CDN providers to have relationships
  - **Compliance:** Where to place traffic controls to keep traffic in defined zones to comply with government regulations

Though a custom approach has strong advantages, this option also carries two significant hazards:

- ***Skillset and experienced staff***: Having the optimal skillset and staff to design, build, and operate a custom cloud is critical. There is no opportunity for trial and error. Work must be done right the first time by experienced staff that has previously successfully built a custom cloud.
- ***Upfront and ongoing cost***: A custom cloud means a SASE vendor is alone financially and requires deep pockets. Not only are there up-front costs to open a custom cloud but ongoing costs will be incurred to expand or refresh. A vendor must be able to develop a balanced business model that allows them to invest to keep customers happy, while at the same time reaching their desired financial metrics.

Netskope is a respected security vendor founded in 2012 that offers numerous security technologies, including CASB, SWG/ZTNA, CPSM, and DLP. The company chose the custom cloud path after achieving unsatisfactory results on public cloud infrastructure. Netskope hired two experienced leaders in 2018, Joe DePalo and Jason Hofmann, to bring carrier-grade design, hyperscale concepts to Netskope and help build the SASE network that they named Netskope NewEdge. Prior to joining Netskope, as Global Head of Internet Services for AWS, Joe was responsible for the global network. Jason was previously Vice President of Architecture at Limelight Networks, one of the largest content delivery network in the world.

In preparing this whitepaper, we had the opportunity to interview Jason. He shared numerous details that showcase three commitments Netskope has made in developing the NewEdge cloud network:

1

Staffing and Technology Investment

- More than 100 dedicated staff were hired before breaking ground on the first data center. Many of these hires have deep protocol and kernel expertise that complements the existing Netskope security team, who have more than a decade of cloud security experience focused on forward and reverse proxy technologies.
- All hardware and software components were purposefully selected, including the latest NVMe SSD, multi-core x86 processors, and 25/100GbE network technologies.
- Over $100 M and counting has been invested in the NewEdge network.

2

Technology Innovation

- The single-pass architecture is optimized for performance with full HTTP/2 support, native interception of TLS 1.3 without down-negotiation, and addresses IPv6 clients via dual-stack.
- The network architecture allows for full compute and traffic processing at every location. All services are available at every data center with no exceptions or additional surcharges.

3

Carrier-Grade Design

- The network is powered by data centers in more than 40 regions globally. Half of these regions were added in 2020, a formidable task during the depths of the pandemic. An additional 12 regions are on track to launch by the end of 2Q21.
- Every region can scale up to 2Tbps or more than 100Tbps globally (Figure 4).
- Netskope maintains full control over its traffic routing, peering, and data center locations, including direct peering at every data center with the Microsoft and Google networks. This represents more than 350 network adjacencies.
- The network service level agreement offers 5-9s (99.999%) of uptime/reliability and a public portal (trust.netskope.com) for real-time service status transparency.

Figure 4: Netskope NewEdge Footprint (May 2021)

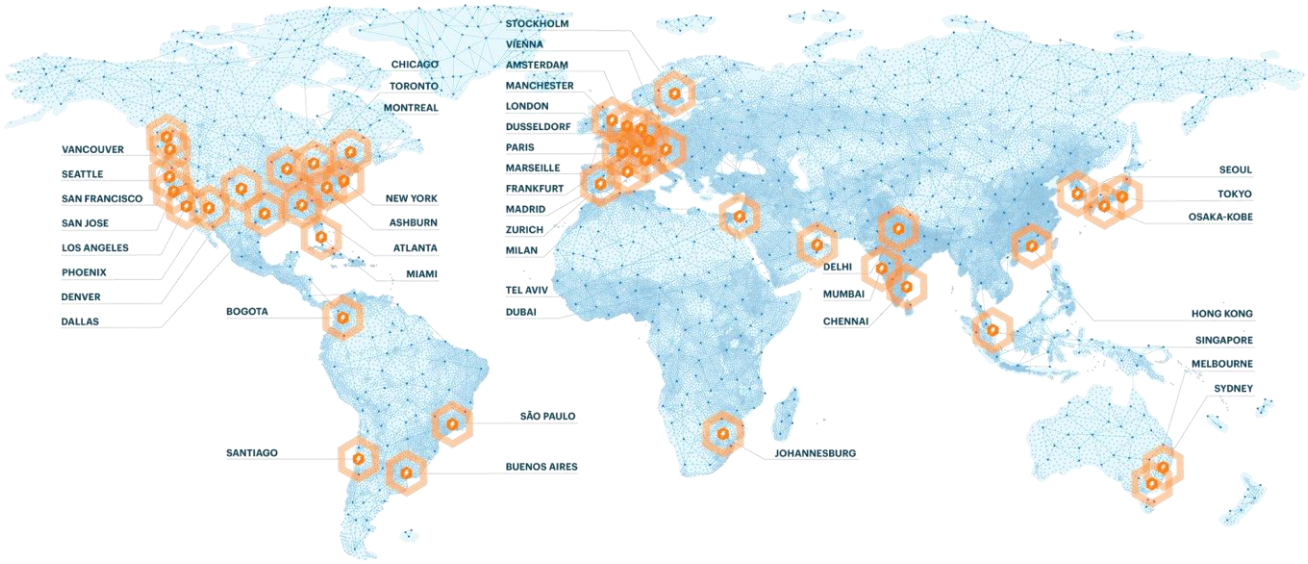

Irrespective of a vendor's claims, the final word should rest with customers. In the last section, we will hear from several organizations we interviewed about the role NewEdge played in their network and security transformation.

CUSTOMER TRANSFORMATION JOURNEYS WITH NETSKOPE

In this section, we delve deeper into the transformation journey that the three organizations we interviewed have undertaken with Netskope:

- Stuart Walters, Chief Information Officer at BDO UK (United Kingdom), the UK member firm with 6,000 employees. BDO UK is part of the $10 billion international network of public accounting, tax, consulting, and business firms operating under the BDO name.
- A senior cybersecurity executive at a global accounting firm with nearly 400,000 employees, who has requested to remain anonymous. This leader is responsible for specifying his company's data-protection strategy across more than 35 global regions.
- Mark Mahovlich, Vice President of Strategy and Execution at ICM Cyber, a cybersecurity consulting firm headquartered in Jackson, Mississippi with 500 clients across 40 U.S. states.

BDO UK

According to Stuart, BDO UK's transformation journey started over four years ago with a new IT direction embracing a cloud- and mobile-first mandate to increase IT automation and digital efficiency. When the journey began, the firm owned on-premises data centers and employed its own IT staff. Later, it moved to a hybrid cloud environment (60% public and 40% private) with 100% outsourced IT staff.

Security is a priority for BDO UK. From the onset of the firm's transformation journey, the BDO team sought to understand unsanctioned cloud services – better known as shadow IT. The team chose to adopt Netskope's CASB solution. In the first phase, desktop agents were deployed and firewall logs were analyzed to understand the extent of unsanctioned cloud usage. Stuart expressed admiration for the depth of Netskope's application awareness and the risk insights provided via Netskope's Cloud Confidence Index that assigns applications one of five risk levels from poor to excellent.

Figure 5: Screenshot of Netskope Cloud Confidence Index (CCI) dashboard and example app drill-down

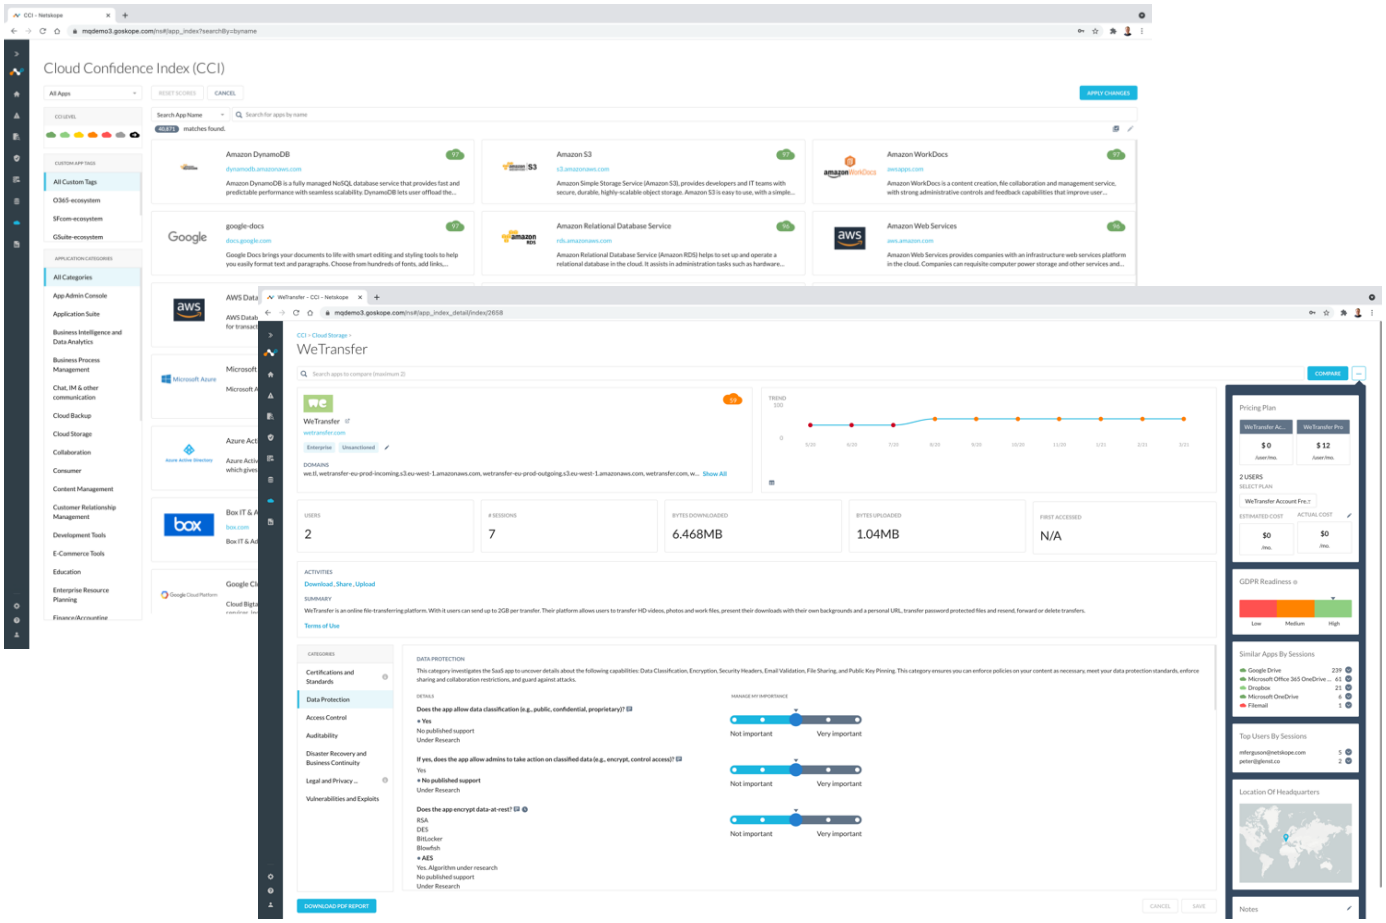

In the second phase, BDO UK advanced from understanding to controlling cloud usage at key network locations. Along the way, the firm worked closely with his clients, the BDO UK departments, and businesses that may have been identified as using unsanctioned cloud services, to find win-win outcomes whenever possible. In the third phase, beginning in mid-2019, the BDO team extended the granularity of control down to individual users by enabling Netskope's cloud-based SWG functionality for inline cloud and web traffic analysis.

Stuart readily admits that the final SWG phase encountered bumps along the way, as is to be expected in any transformative change. Firewall rules and email configuration had to change. Some services that previously worked were temporarily broken. But BDO UK persevered and eventually was successful in smoothing out all the bumps by the beginning of 2020 – fortuitously only weeks ahead of the impending COVID-19 pandemic.

Prior to the pandemic, few BDO UK employees worked remotely. But one fateful Monday afternoon in mid-March 2020, the pandemic hit, and the mandate to work from home came down. BDO UK's IT infrastructure was put to the ultimate test. On the very next day, all 6,000 BDO UK employees logged into work from home for the first time. They did so seamlessly without any hiccups. For this success, Stuart credits the combination of his strong cloud-first application environment, providing the necessary elasticity, with the mobile-first strategy leveraging Netskope's cloud-based SWG on NewEdge, which offered the crucial connectivity substrate for users to connect to their applications without faltering. The work anywhere era had arrived at BDO UK with Netskope NewEdge as a key element in its achievement.

Global Accounting Firm

Whereas Stuart gave equal priority to the twin goals of IT modernization and increasing data security, our second end-user, the senior cybersecurity executive at a global accounting firm we previously introduced, placed utmost priority on preventing data leakage. In his words, "Keep the bad stuff out and the good stuff in." He had already deployed a client-based DLP solution and wanted to achieve defense-in-depth by adding a network-based DLP solution based on CASB and SWG technologies.

His security team partnered with the network team to ensure that any new solution they chose would support the network in two important ways:

- ***Support network exceptions, as needed:*** The new solution had to be flexible and provide a good fit in the existing network environment. For example, VPNs with split tunnels needed to continue to function for legacy applications. For a long time, the organization had embraced a cloud-first mentality for all applications. Yet, a minority share of legacy applications still necessitated exceptions.

- ***No change in end users' application experience:*** The expectation was that the new solution would maintain end users' application experience and not introduce unacceptable performance throttling or latency.

Based on Netskope's strong CASB reputation, the customer chose to deploy this solution more than one year ago. At present, both the security and network teams continue to be pleased with Netskope. In the customer's view, Netskope delivered a stable product and he praised the network's flexibility to set security and network policy exceptions. The commendation continued when we discussed the end-user application experience. Not only was Netskope capable of processing terabytes of cloud and web traffic but it had also improved the throughput of certain applications by up to 50%.

ICM Cyber

While global enterprises have senior thought leaders, like Stuart at BDO UK, to internally drive IT transformation, many enterprises – especially small to medium enterprises – rely on external help. ICM Cyber fills this external consultancy role with folks like Mark at the helm.

All of ICM Cyber's clients seek to address one or more pain points associated with legacy architecture. Yet only about 30% arrive at ICM Cyber knowing the solution they need. This means that in 70% of cases, ICM Cyber helps to navigate the uncertainty.

Usually, one of three IT teams comes to ICM Cyber for help: the security operations team, the application team, or the networking team. However, Mark notes that in the last few years, the silos between teams are breaking down and jointly engaged teams are increasingly approaching ICM. The most common pairing is the security and networking teams, as network teams finally realize the need to embrace security.

We previously described that network teams want to embrace transformation. Yet they are hesitant out of concern they will introduce a higher risk of network failures.

Mark explained the three ways he appeases a network team's concerns:

He selects a cloud-based toolset that provides equivalent end-to-end visibility, metrics, and diagnostics like the network team currently uses.	*He trains the network team on how the cloud-based toolset integrates with their standing network configuration and troubleshooting workflows.*	*He launches a pilot program that runs a production application just as well, if not better, on the toolset versus in the legacy environment.*

Mark's preferred cloud-based toolset is the Netskope platform. For Mark, Netskope embodies the longstanding mindset at ICM Cyber that organizations need to deploy converged, cloud-based solutions for the sake of productivity and actionable intelligence. Managing dozens of point solutions is no longer tenable, according to Mark. For him, Netskope is the path forward to enable a highly mobile, cloud-driven, work-from-anywhere enterprise.

CONCLUSION

As we enter the post-pandemic era, enterprise networking and security strategy must change to address the forthcoming new normal of work from anywhere. Traditional approaches that assume clear network perimeters exist or rely on legacy MPLS and VPN technologies will no longer succeed. Enterprises need to embrace a cloud- and mobile-first IT strategy in the manner of the organizations we interviewed. In the new normal, networking and security are intertwined and collectively addressed using a SASE-based approach.

For many enterprise teams, the shift toward SASE is daunting. In particular, network teams want to embrace the change but they are extremely reticent out of concern for having a negative impact on their networks. But the path forward need not be fraught with failures if network teams select the right tools and take the stepwise approach described by several of the leaders we interviewed.

Netskope was the right security platform for those we spoke with due to its strong CASB, DLP, and SWG security capabilities built on top of a resilient cloud network – the Netskope NewEdge Network. All of our interviewees expressed satisfaction with the NewEdge network, whether due to its robust performance to keep end-users happy – such as when BDO UK took its 6,000 users remote in a single day without hiccups – or the end-to-end visibility and diagnostic capabilities that enable network teams to perform their role.

As enterprises evaluate their options and decide which SASE journey to take, we hope our paper sheds new light, especially for network teams, on an increasingly traveled path. We believe vendors that build, converge, and improve upon standalone SD-WAN and SWG solutions, such as Netskope, should be on the consideration list for an enterprise seeking to address the needs of the new work-from-anywhere era, while also simplifying, modernizing, and optimizing their enterprise network.

About Author

Mauricio Sanchez joined Dell'Oro Group in 2020, and is responsible for Network Security & Data Center Appliance market research program, as well as Security Access and Service Edge (SASE) Advanced Research Report. He brings over 20 years of experience as an executive manager in networking and security technologies, products, and solutions spanning data center, campus, and mobile architectures. Mr. Sanchez helps shape the coverage of next-generation networking architectures and services models. Mr. Sanchez's research and analysis has been widely cited in leading trade and business publications. Mr. Sanchez is a frequent speaker at industry conferences and events.

Email: mauricio@delloro.com

About Dell'Oro Group

Founded in 1995 with headquarters in the heart of Silicon Valley, Dell'Oro Group is an independent market research firm that specializes in strategic competitive analysis in the telecommunications, networks, and data center IT markets. Our firm provides world-class market information with in-depth quantitative data and qualitative analysis to facilitate critical, fact-based business decisions. Visit us at www.delloro.com.

About Dell'Oro Group Research

To effectively make strategic decisions about the future of your firm, you need more than a qualitative discussion – you also need data that accurately shows the direction of market movement. As such, Dell'Oro Group provides detailed quantitative information on revenues, port and/or unit shipments, and average selling prices – in-depth market information to enable you to keep abreast of current market conditions and take advantage of future market trends. Visit us at www.delloro.com/market-research.

Dell'Oro Group

230 Redwood Shore Parkway
Redwood City, CA 94605 USA
Tel: +1 650.622.9400
Email: dgsales@delloro.com
www.delloro.com