SANS

A SANS Whitepaper

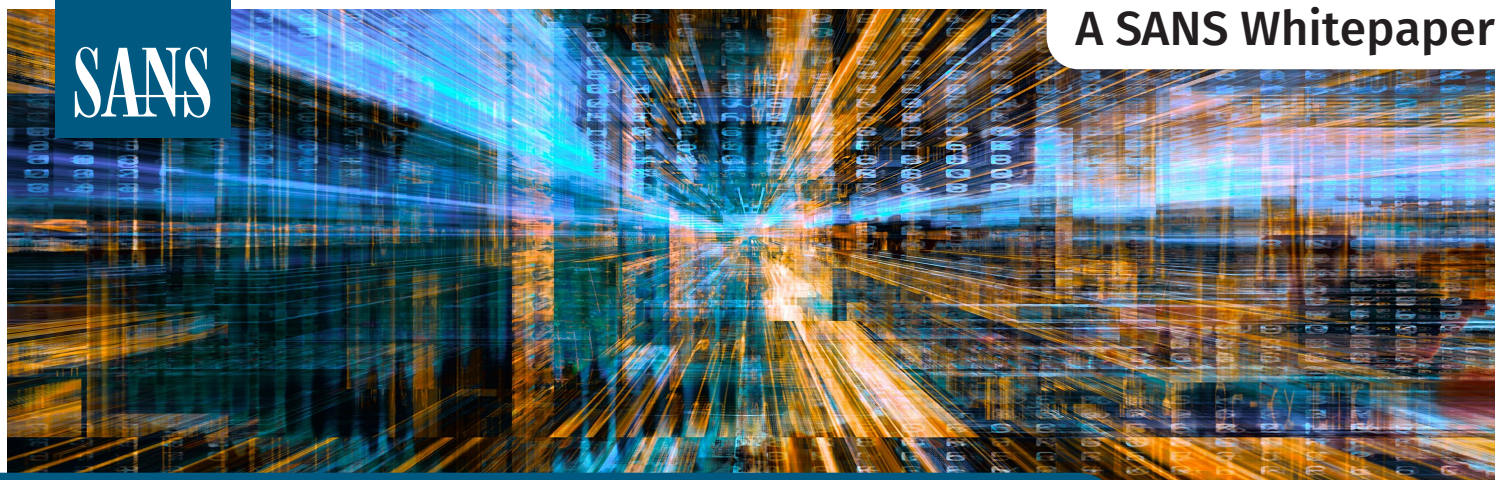

How to Manage the Shift to Cloud Security

Written by **Dave Shackleford**

December 2020

Sponsored by:

Netskope

The Shift to Cloud Security

More and more organizations are moving data, systems and applications into cloud environments. Some of the reasons for this shift include:

- **Cost savings—**Many organizations can save on legacy costs of building, operating and maintaining on-premises application environments and networking by moving to cloud-based services.
- **Operational efficiency—**Without the need to patch, update and operate many elements of the IT infrastructure, providing direct-to-cloud access, organizations can better focus on building the applications and services necessary to business operations.
- **Improved speed and scalability of deployments—**With the advent of DevOps and software-defined infrastructure, more organizations are building and deploying workloads and applications more rapidly than ever, while reducing manual errors and tasks.
- **New features and technical capabilities—**Cloud-based services are often leading in innovation, providing more cutting-edge capabilities than organizations can possibly build themselves on-premises.
- **Remote workforces and business continuity—**The COVID-19 global pandemic abruptly placed workers outside of their corporate offices, requiring IT teams to be agile in connecting their users to the applications and workloads they need to keep businesses operational and effective.

Figure 1 is only a snapshot of several key business drivers motivating organizations to adopt SaaS and move workloads and applications to the public cloud. Meanwhile, security teams are still working to improve security capabilities and effectiveness in cloud deployments and applications. Because cloud usage has grown significantly, it's critical to provide a wide variety of security controls, ranging from cloud workload attack monitoring and control to protecting the users who access cloud applications, including managed SaaS, private applications and shadow IT. In the Verizon Data Breach Investigations Report for 2020, cloud assets were involved in 24% of breaches. In 73% of those breaches, the attacks involved an email or web application server, and 77% of these breaches also involved compromised credentials.[1]

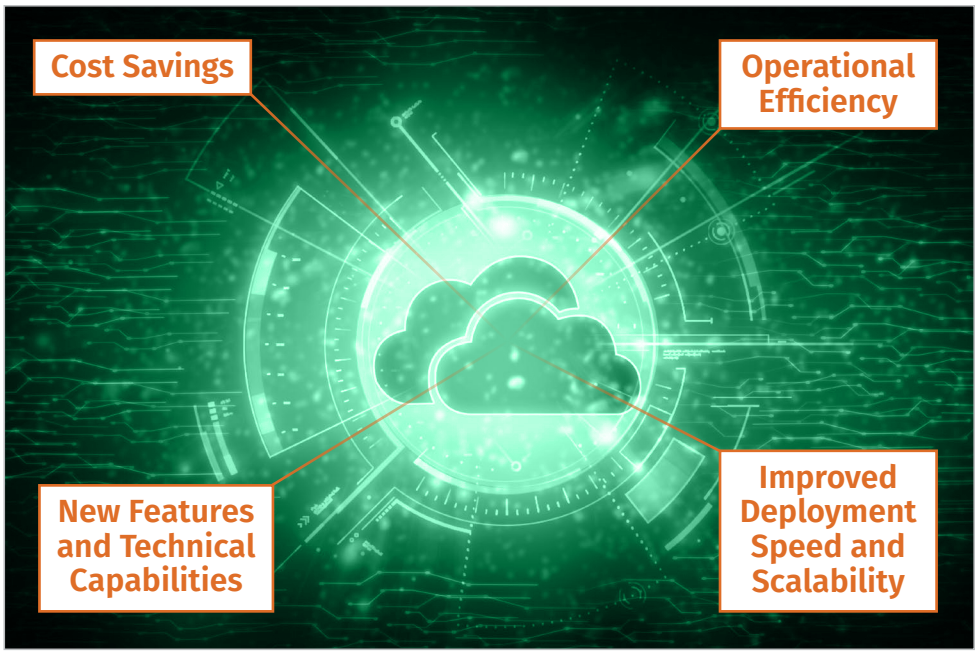

Figure 1. Reasons for the Move to Cloud Security

This trend is compounded by a significant shift toward remote work and "untethered" access, where end users, developers, IT operations and engineering, DevOps and other users can access cloud applications and resources (of all types) directly from their own systems in any location. Increasingly, organizations seek the right combination of enterprise collaboration tools to handle an entirely remote workforce. They want to implement the best practices to secure users who need access to a wide variety of cloud applications and environments, and they want to enhance security controls and capabilities for IT users and DevOps teams who need to continue deploying assets into the cloud. In 2020, we saw a rapid shift into this work modality due to the global COVID-19 pandemic. In Netskope's 2020 Cloud and Threat Report, roughly 30% of many organizations' workforce were remote in March 2020 when the pandemic was announced; by the end of April, roughly 64% of employees were working remotely.[2] See Figure 2 for a graphical depiction of this trend.

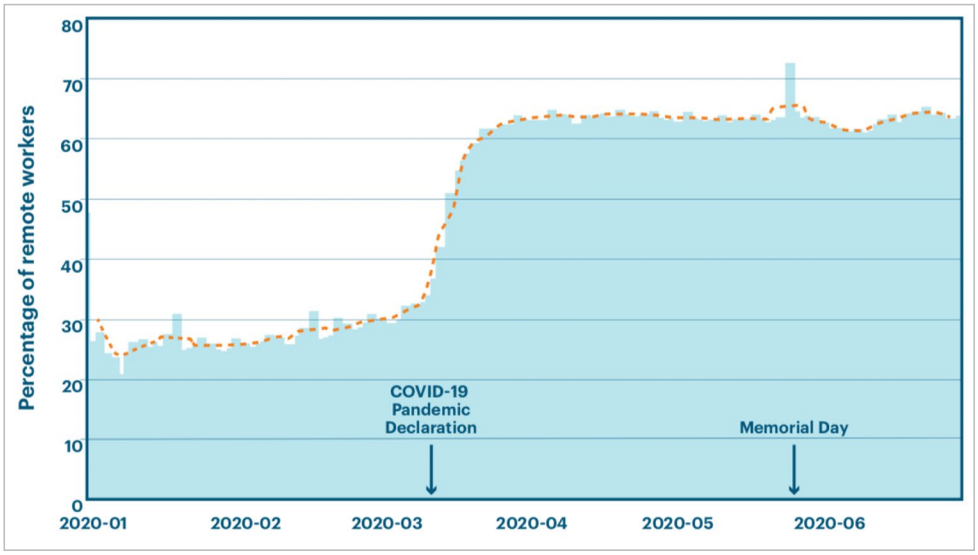

Figure 2. Trend in Remote Workforce in 2020

[1] "2020 Data Breach Investigations Report," https://enterprise.verizon.com/resources/reports/dbir

[2] "Cloud and Threat Report, August 2020 Edition," www.netskope.com/lp/cloud-and-threat-report/?utm_source=google&utm_campaign=NA-Brand&utm_agn=&utm_content=419914836918&utm_term=+netskope&gclid=Cj0KCQjwreT8BRDTARIsAJLI0KKqT2R9qOMQ8aSX_hIYFfMoXVFuCg3YyvD3xFh2mobyJ4tJyTQoMtgaAkQDEALw_wcB [Registration required.]

The Netskope report indicates that more applications are in use, too. Large organizations of 2,000–4,000 users were accessing roughly 1,148 applications; larger organizations accessed almost 3,000! In addition, there was a significant change in some users' behavior, with:

- A 161% increase in visits to high-risk apps and sites
- A 600% increase in visits to adult web content
- An 80% increase in the use of collaboration tools and applications
- 63% of malware coming from cloud storage

The use cases for cloud applications and services continue to expand rapidly, and organizations are realizing that many types of access scenarios are shifting as we move toward more SaaS use, hybrid cloud infrastructure deployments and multicloud deployment and interconnectivity. Whether oriented toward end-user access to applications and services or traditional data center access to branch offices and other remote locations, the need to make our traditional data centers the hub of connectivity is rapidly becoming more of a hindrance than a help. Accordingly, we've seen the emergence of a new service model named by Gartner as secure access service edge (SASE), which combines a number of different elements of cloud services, networking and security into a unified fabric.

What Is SASE and How Can It Help?

With the need for more software-defined infrastructure accommodation and end-user cloud service access, SASE has brought a number of cloud security and infrastructure enablement technologies together into a unified platform that combines the following elements:

- Software-defined wide area network (SD-WAN)
- Secure web gateway (SWG)
- Cloud access security broker (CASB) (both inline and API-based access to SaaS)
- Zero trust network access (ZTNA)
- Firewall-as-a-service (FWaaS)
- Cloud data protection

A number of other features are becoming more common, too, such as:

- DNS protection
- Web application and API protection (WAAP)
- Remote browser isolation
- Network malware sandboxing
- VPN replacement (managed and unmanaged devices, access controls)

In order to qualify as a cloud-native SASE, the technology must be software-based and totally hardware-neutral, with the same elastic capabilities to scale as any cloud gateway solution available today. Given the focus on broad coverage and connectivity, a SASE must offer numerous globally distributed points of presence (POPs) to connect through for office locations and end users alike. Because these services will be processing and evaluating data and content for potential exposure and/or threats, SASE must offer scalable encryption and decryption with "single-pass" scanning for data and threat protection. Single-pass scanning allows for minimal disruption of user activity while still offering critical security services at scale. Compare this with services like email that have more tolerance for "queued" scanning that can take somewhat longer. SASE services should be designed for multitenant access and not loosely coupled together from acquisitions (a common strategy for service providers that are strong in one category initially, such as SD-WAN or SWG). Figure 3 presents these qualifications in one graphic.

Secure access server edge (SASE) combines several cloud security and infrastructure enablement technologies into a unified platform that is intended to deliver to the user the set of network and security services when and where needed for their use cases.

SASE brings about a wide variety of changes in how organizations and their users access computing and application resources. The first element of SASE is oriented toward network access, control and architecture. Software-defined networking and security has come to include SD-WAN, which allows for interconnectivity between on-premises environments and cloud provider infrastructure through a singular backbone service or vendor solution. These networking services often provide common networking capabilities such as routing, bandwidth shaping and quality of service (QoS), and core content delivery network (CDN) services that can set priorities on specific content and service access and transmission. As with any newer service architecture, there is some confusion about what the key differences are between this and more traditional networking architecture and controls. Table 1 on the next page should help enterprise networking and security teams to gain a better understanding of the differences.

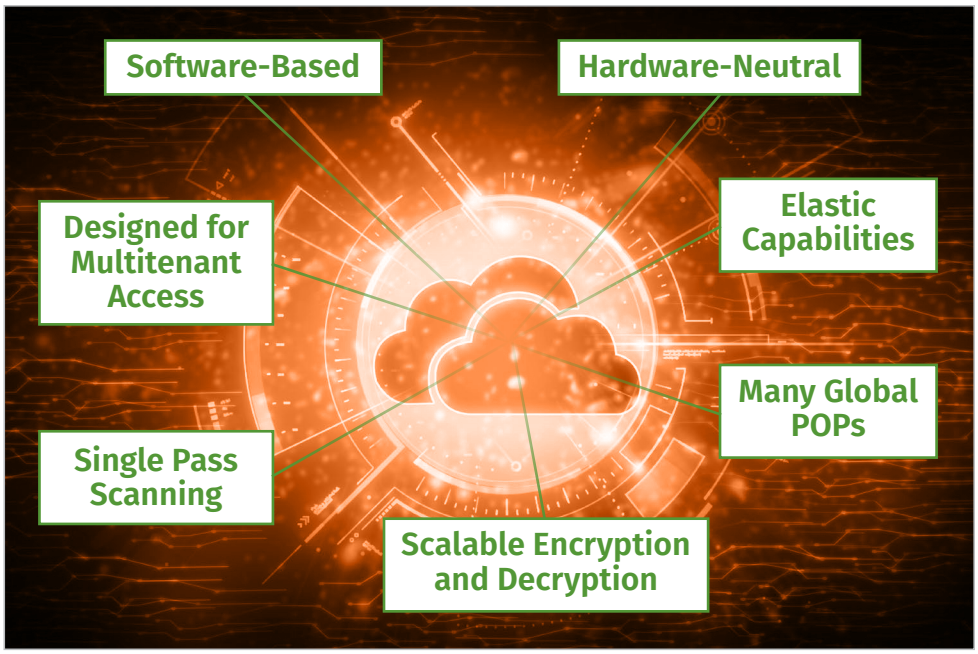

Figure 3. Cloud-Native SASE Qualifications

Table 1. Differences Between SASE and Traditional Networking Architecture and Controls

Features/Controls	Traditional Models	SASE Models
Remote access to on-premises resources	Most traditional models have largely relied on VPN technology through SSL/TLS browser access or a dedicated endpoint client.	SASE services could largely act as a VPN replacement, where users connect to the SASE service for access to both on-premises resources and cloud services. All policy is defined and applied through the SASE console using zero-trust network access (ZTNA).
Access to cloud resources	On-premises network access to cloud resources treats these like any other online properties, using traditional firewalls, proxies and routing controls.	SASE services provide optimized and streamlined network access that is more "cloud aware" for numerous SaaS, PaaS and IaaS provider offerings. This may rely on API integration and request introspection for end-user requests, but most implementations are inline through forward proxying and agent-based deployments.
Network access controls	Most on-premises environments rely on switching, routing, firewalls and proxies for access control to resources.	SASE services aggregate a number of network security and access controls (including FWaaS, forward and reverse proxies and others) into one unified fabric.
SD-WAN, WAN optimization and bandwidth aggregation	All of these controls and capabilities usually require several different vendors and products to function, and they may lack in integration.	A SASE service integrates SD-WAN access and traffic optimization capabilities into a single gateway service for all access types.
SWG	SWGs are usually separate proxy appliances focused on web traffic and SSL inspection with migration to cloud-hosted SWGs going forward.	SASE platforms integrate SWG policies and services into the same single-pass approach alongside inline cloud inspection for apps and cloud services.
Threat detection	In a traditional on-premises model, threat detection is accomplished using next-generation firewall (NGFW) platforms, network malware detection sandboxes or endpoint detection.	SASE services combine numerous network threat detection capabilities into one service fabric. Examples include sandboxing, machine learning, remote browser isolation (RBI) and detecting fake phishing forms and rogue app instances.

Here is a partial list of SASE use cases that may be most appealing or interesting to organizations rapidly shifting to use more cloud services:

- **Core network—**For internal users and services that need to access cloud services, SASE may replace independent CASB and SWG solutions and other independent options to access the cloud securely.
- **IoT/edge—**IoT and edge network scenarios can be isolated and connected through network-restricted policies assigned within a SASE service edge or in tandem with SD-WAN vendors.
- **Branch—**Branch locations can connect to on-premises and cloud resources through SD-WAN and SASE service edge POPs instead of traditional network service providers using MPLS to backhaul traffic.
- **Remote end users—**Remote end users will now connect to SASE services to receive access policies traditionally applied by internal VPN solutions. In essence, SASE offers a "VPN as a service" that can secure connections to any approved destination using a zero-trust model.

ZTNA

One of the more intriguing capabilities offered with SASE is identity-driven access management. The first major shift in the way a SASE approaches access management is the definition of what constitutes an identity in the first place. Although the more traditional concept of identity still applies (users, groups and role assignments), all edge locations and distributed WAN branches and network origins are considered identities, as well. In a cloud-focused enterprise, secure access decisions should be centered around the identity of the entity at the source of the connection (users, devices, branch offices, IoT devices, edge computing locations and the like).

The identity of the users, groups, devices and services in use remains the primary element of SASE identity access policies. SASE identity policies are evolving, however, to include a number of additional relevant sources of identity context that can factor into policy decisions and application. These sources may include some combination of the identity's location (using geolocation tracking and tagging), the time of day, device security evaluation or trust validation, and the sensitivity of applications and data entities are trying to access. These factors can help organizations to develop and refine a more progressive least-privileged access strategy that enables strictly enforced access control. The promise of SASE identity policies is that organizations will be able to control interactions with resources based on more varied relevant attributes including application access, entity identity and the sensitivity of the data being accessed.

This is a shift in the security landscape that has been underway for some time in the form of ZTNA and microsegmentation based on applications and identity affinity policies. To date, this has been a largely internal technology shift but has now branched out to a broad access control methodology that facilitates identity-based controls for entire office locations, remote users, IoT devices and more. This model looks to significantly improve upon the classic access strategies that focus only on network information like IP addresses and ranges, or network edge devices with rigid connection methods that may be complex to set up and maintain. This shift to policies oriented toward application, data, device, and user affinity policies may streamline the creation and management of access policy management. Once authenticated and authorized to access resources, a SASE service can then act as a VPN-like broker that protects the entire entity session, regardless of where it originates and connects to. In keeping with the theme of zero trust, SASE solutions should have flexible options to apply end-to-end encryption of sessions that can also layer in additional cloud and web application protection, API inspection and security assessment, content inspection for DLP, UEBA for insider threats and any other variety of security services in a single-pass gateway access model.

The promise of SASE identity policies is that organizations will be able to control interactions with resources based on more varied relevant attributes including application access, entity identity and the sensitivity of the data being accessed.

Many types of attacks are likely to be mitigated with effective application of SASE services in the future. Some of these are simple. In the networking realm, with strong unified policy management, more thorough validation of branch office connections, approved IoT devices, and edge services and locations can be built and maintained. SASE should help to curtail some man-in-the-middle interception attacks, spoofing scenarios and malicious traffic of some types. End users should also benefit from this model, because leading SASE providers can help to securely encrypt all traffic from remote devices regardless of location and even apply more rigorous inspection policies based on public access such as airports and coffee shops. Depending on the identity of the user and originating device, privacy controls can be better enforced by routing traffic POPs in specific regions, as well.

The move to building access models around identity will take time and likely some substantial level of effort at first as we move away from tiered access models based on IP addresses. However, the end results will likely be worth it, making security operations more efficient and attacks more difficult for adversaries at the same time.

Additional Security Services and Features

At the core of SASE is a dual emphasis on user and data security (with network security as another key shift). This emphasis can help to enable data and threat protection for both cloud and traditional web access. Cloud security-as-a-service (SecaaS) is the final convergence category that is included in SASE, and it helps add additional data security and end-user security controls and capabilities that a well-rounded SASE service should provide. SecaaS is a very broad category, layering in services often provided by CASBs such as DLP, content filtering, malware detection and response, cloud provider reputation scoring, user behavioral monitoring and more. In addition, SASE brings together more SecaaS offerings such as SWGs, network intrusion detection and prevention, and even remote browser isolation (RBI). This newly defined cloud networking and security category will definitely lead to a shift in some cloud security service providers changing and updating their offerings to include features not offered before. In essence, the SASE space looks to take advantage of the cloud brokering model already seen with CASB, CDN, and even identity-as-a-service (IDaaS) to include more networking capabilities and control, as well as combined security services in one single-pass gateway model that could potentially simplify networking and security control stacks.

SASE offers some distinct opportunities. A single provider could offer cloud service and internet access to end users, data center services and platforms, and IoT and other distinct devices through a combined networking and security fabric that is jointly defined and administered by networking and security teams (likely with input from mobile, application development and systems administration teams as well).

Features to Look for in SASE

When evaluating SASE services, consider the following services and capabilities:

- **Broad applicability to users/devices of all types/locations—**All types of end users and remote devices should be accommodated, with different policy application depending on the location and type of user.
- **Cloud-centric focus for SaaS and user/data security (identity/data as a perimeter)—**As more end-user services shift to the cloud, there should be more direct integration of inline CASB (forward proxy for managed devices and reverse proxy for unmanaged devices) with SWG (forward proxy for managed devices) in a single-pass solution.
- **Data protection (granular policy controls by app and instance) with DLP—** Data protection and monitoring controls are usually associated with CASB providers, and SASE providers should have these same detection and control capabilities that are real time (data in motion) and via API (data at rest).
- **Threat intelligence and visibility into SaaS, IaaS/PaaS, and web content/context—** A sound SASE solution should have broad visibility into attacks occurring across numerous customer scenarios, and this should be distilled into threat intelligence for security operations teams.
- **Use of machine learning (ML) for advanced threats, data classifiers for DLP, and behavior anomalies—**An advanced SASE platform should have advanced machine learning capabilities with deep security analytics to accurately process and report on detected incidents.

With all of these capabilities in place, organizations can turn to the types of use cases that may benefit from SASE.

SASE Use Cases

SASE can help organizations to improve connectivity and cloud security scenarios in several use cases, including the following:

- **Spot unintended/unapproved data movement.** SASE could potentially help to identify sensitive data accessed or migrated into/out of the cloud and between company and personal application instances, or to third parties.
- **Control user activity from managed and unmanaged devices across all apps, including personal instances.** As more organizations have had to allow users to work remotely from unmanaged devices, the need to implement security controls has moved from data center appliances and back-hauling traffic to a cloud security service-edge fabric.

- **Protect enterprise data and intellectual property with data protection and advanced DLP.** Deep data protection capabilities are common drivers for SASE. This need often follows CASB implementations, where organizations are concerned about data loss/exposure in apps and cloud service scenarios including shadow IT apps.
- **Detect abnormal behaviors using machine learning anomalies for users, apps, and data-in-motion for cloud and web traffic.** With deep security analytics, SASE providers can help to rapidly identify unusual patterns of behavior for users with compromised credentials, accidental data exposure, or insiders with intent to exfiltrate data.
- **Impart conditional/contextual access to cloud/web services.** One of the more important controls SASE can provide is an understanding of risk for apps, user confidence, device type, data classification and activity for conditional and contextual access.
- **Provide granular cloud app controls, selectively enable apps to third parties and exceed compliance regulations.** CASB controls that protect users from accessing and interacting with specific cloud applications or functions are critical and also align with consolidated SWG and CASB controls for a single-pass SASE architecture.
- **Identify and block cloud-enabled web and SaaS threats.** More malware than ever (including ransomware and other highly malicious code) is being accessed in and disseminated by apps and cloud services. SASE platforms should be able to identify and block a wide variety of threats including cloud phishing, cloud-hosted malware, fake forms for access compromise and cloud-based callback communications (C2). This feature should also include protection from threats in apps, browsers and embedded links, often cloud-enabled stages using trusted domains and valid certificates.
- **Deliver broad and efficient application access.** With more users accessing applications from anywhere, secure and fast access is imperative. Speed and efficiency for round-trip time (RTT) is definitely a primary focal point of SASE services.

Security teams can explore many different avenues with SASE, given the convergence of different security controls and capabilities. Where, though, should organizations get started?

How to Get Started with SASE

Organizations considering SASE providers based on networking and networking security capabilities should take several key items into account. First, heavily scrutinize uptime and availability SLAs, along with the breadth and types of POPs for connectivity. Second, carefully assess network and network security capabilities. Very few SASEs are good at all network functions and services collectively, so evaluate their strengths and weaknesses—some may excel at SD-WAN and traffic optimization, whereas others are more focused on CASB, SWG, and/or ZTNA, for example.

For most organizations, the initial focus for SASE will be on next-generation SWG combining inline CASB with single-pass SWG security controls, primarily for enhanced cloud access visibility and control. More sophisticated solutions will share threat intelligence and indicators of compromise (IoCs) with endpoints and SIEMs, and integrate with identity and access management (IAM) for more flexible access controls including multifactor authentication (MFA) and step-up authentication/authorization based on risk policy controls. For more advanced organizations and scenarios, ZTNA for access control to private applications likely makes sense as a starting point.

Conclusion

For organizations considering whether SASE options may be a good fit, there are some key considerations. First, decide whether a unified strategy with a single provider for numerous critical services makes sense to you. The primary benefit would be operational simplification and a smaller list of vendors and providers; the trade-off may be a massive single point of failure or exposure. Second, carefully scrutinize capabilities. Most SASE vendors started as something else, such CASB, SWG, or SD-WAN service vendors, and are now "bolting on" other capabilities through acquisition or scrambling to develop them quickly (sometimes with mixed results). Lastly, cost, both operational and financial, is always a major factor in decisions like this, so you'll need to factor that into your decision-making process.

SASE is a new category that likely has some significant maturing to do, as well. If you have most or all of the capabilities you need today, then understand your use cases and the compelling reasons to change—this service fabric convergence is occurring naturally in the cloud space and will come together organically.

About the Author

Dave Shackleford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Sponsor

SANS would like to thank this paper's sponsor: