

8 Key Considerations for Networkers When Assessing Security Service Edge Platforms

Before the cloud, enterprise network architects focused on securing connectivity to and within the data center perimeter. Today, we're entering a cloud-first era of enterprise networks and applications where data is no longer contained within the data center perimeter. Rather, it resides outside of the organization, distributed among Software-as-a-Service (SaaS) and cloud environments.

Network architects are now focused on providing users with secure, fast, and reliable access to increasingly dispersed applications and data. With many users now working remotely, accessing applications and data in diverse locations, providing this secure access has become much more challenging. As a result, network and security teams are coming together to adopt a Secure Access Service Edge (SASE) architecture powered by a Security Service Edge (SSE) platform.

In this eBook, we provide recommendations to network architects based on our experience working with infrastructure and security teams in enterprises to vet SSE solutions. This eBook will also provide insights into the choices we made when creating our own security cloud and the Netskope Intelligent SSE solution, specifically addressing components providing secure access to web, cloud, and private apps.




8 Key Considerations for Networkers When Assessing Security Service Edge Platforms

| | |
|--|----|
| Section 1 The requirements of an SSE Platform | 04 |
| Section 2 The placement of processing | 05 |
| Section 3 Management and visibility | 07 |
| Section 4 Deployment options | 08 |
| Public compute/storage cloud | 08 |
| Carrier network cloud | 09 |
| Content delivery network | 10 |
| Purpose-built private cloud | 11 |
| Section 5 First and last mile latency | 13 |
| Section 6 Processing latency | 15 |
| Section 7 WAN interconnect | 16 |
| Section 8 Resiliency | 17 |
| Summary | 19 |



The requirements of an SSE platform

The first step for network infrastructure architects is to understand the requirements of an SSE platform. Broadly speaking, it must:

- 
- 1.** Connect enterprise users regardless of location, network, or the device(s) they happen to be using.
 - 2.** Ensure fast and reliable access to enterprise applications with zero trust principles in mind (so that users aren't tempted to bypass the security cloud's controls).
 - 3.** Control access to applications, no matter where they reside.
 - 4.** Inspect traffic flowing between users and cloud applications to enforce relevant policy controls with contextual and instance awareness.
 - 5.** Maintain proper localization and proximity information for each user/service combination.
 - 6.** Work together with the information security teams to integrate security by design, such that security controls are weaved into every aspect of network connectivity.
 - 7.** Consider business governance, compliance, and data protection requirements.
 - 8.** Not compromise any aspect of performance, availability, or security in its delivery model

The key requirements for an SSE platform can therefore be summarized as a distributed implementation with relatively intense computations performed at a relatively large number of locations within proximity to all users, peered with cloud and SaaS providers, combined with the management and visibility tools required to deliver an on-demand, elastic, and high-performance interconnected security service.

The placement of processing

With the requirements of an SSE platform mapped, architects should next consider where best to place processing tasks in order to minimize sources of friction that may adversely affect application performance. In the past, network architects usually had to deal with trade-offs between availability, performance, and security. To properly secure data flows, traffic needed to pass through a security service chain that consisted of various appliances, with each inspection type introducing more latency and degrading performance. The challenge has always been a struggle, as providing one of these requirements impacts the others. In this regard, there are **four main factors to consider**:



1. The distance (measured by the time required to send data) from the platform to the end user.
2. The cost of security processing (inclusive of resources, time constraints, and latency).
3. The distance of security processing to the service or application.
4. The type and depth of security inspection.

As any network architect knows, good application performance depends on short distances between services using efficient routing. The aim is to avoid diverting traffic to a central location, which can become a choke point, or geographically dispersed network locations, which introduce latency due to distance. It also requires minimal latency when performing any traffic inspection. Performance considerations therefore demand that architects put security services near to their end users. An effective Security Service Edge platform is one where highly efficient processing takes place at multiple locations around the world, or to put it another way, where there's a lot of compute in a lot of places.

A key consideration when appraising vendors is to go beyond raw data center, zones, or virtual points of presence counts and ascertain the number of locations or regions where processing occurs in close proximity to your end user, and that each processing location is capable of running all services inline.

As of October 2023, **Netskope NewEdge** is powered by full compute data centers in

70+ regions worldwide.

Always ascertain how many of a vendor's claimed zones/regions/data centers are simply "front doors" or virtual points of presence (POPs) that do not offer full compute or inline inspection.



As of October 2023, Netskope has enhanced its region coverage with an additional number of "**Localization Zones**" that offer an egress IP in

200+ further regions

to increase resilience and enhance the digital experience, so users get content and apps tailored to their location, plus they can always access their critical apps, even those that are geo-fenced.

Management and visibility

The second major architectural consideration focused on visibility and management. Network infrastructure architects need a clear and consistent view of the operating characteristics of the network access and data flows, as well as the security cloud's connectivity and processing resources. Architects should strive for a comprehensive operational status and health visibility of users' traffic based on real traffic flows as well as synthetic telemetry data. In addition, architects need a single console for easy deployment, management, and troubleshooting, through which they can adjust the cloud implementation as required. Critical capabilities include capacity management tools to determine the levels of utilization, view of latency metrics, service consumption, geo-specific service usage, user counts, private application usage and performance metrics, site and tunnel health, and overall performance visibility.

In some ways this approach runs counter to the usual cloud story, where IT teams are largely happy to leave visibility and control in the hands of the cloud service provider. However, when it comes to a Security Service Edge platform, the ability to intercept, decrypt, inspect, and re-encrypt connections without impacting performance is fundamentally important. Visibility and control are non-negotiables and a requirement for operations teams to properly troubleshoot user connectivity and performance issues. The goal should be to simplify and accelerate operations through the use of a single, unified management console and enforcement point.



Deployment options

With the fundamentals addressed, we can take a look at how best to technically realize an SSE platform. There are four core options open for consideration.

Public compute/storage cloud

One simple approach that vendors may choose is to leverage one of the major public clouds such as AWS, Azure, or Google Cloud Platform (GCP), to name just three. The benefits of doing so include:



- Immediate access to multiple locations for security processing.
- These locations are accessible through multiple networks.
- Architects can focus on tool integration, security scanning, and policy enforcement, as the underlying hardware considerations are all taken care of by the hyperscaler.
- Rapid time to market.

However, there are also several drawbacks to this option, which means this approach falls short of what's needed for an effective network access and security cloud.



- Although cloud providers have sites around the globe, they limit you to regions that may not be where your users or applications reside.
- You may choose a vendor whose platform is built in the same environment as your own applications, but how do you enable seamless, efficient access if your applications are not serviced by that provider, or you move to a multi-cloud architecture?


One of the most significant factors is that the cloud architecture has been built to support the hyperscaler's business model, not those of its customers. As a result, its computing and networking resources are distributed to support its aggregated customer base and cost points. Security processing resources will therefore often not be close enough to a business's users and applications to deliver the required performance.

Carrier network cloud

Another design option is to use a large public data network. Networks of this kind offer a vast number of locations and connect to an enormous number of users, providing the broad reach required for an effective access and security cloud. In this approach, the carrier is a one-stop shop, providing all the transit and peering capabilities needed.

Once again, however, there are a number of drawbacks. First, the platform has no visibility beyond the traffic that's going in and out of the carrier's facilities. Typically, a carrier's traffic engineering and capacity management tools are not shared with customers, so architects do not have the visibility they need for an effective platform.

Another issue is performance. A single carrier can work well at moving data between its own network locations, but that's not the crucial problem to solve when building an SSE platform. The platform needs to solve a version of the "internet performance problem," and networking experts have known for a long time that isn't possible with just one network operator.



Content delivery network

One method that has been used for decades to improve network performance is content delivery networking (CDN). A CDN uses geographically distributed and interconnected servers to provide content to a user. They provide cached internet content from a network location closest to a user.

In essence a “virtual network,” a CDN uses the networks that form the internet selectively to achieve the best possible performance. CDNs appear to be well-suited to host an SSE platform, as they are present across diverse networks, and they have edge computing capabilities that could conceivably be used for security processing.

In reality, however, the promise of CDNs falls short. While the best use hundreds of thousands of servers across thousands of locations, these resources are also shared between thousands of customers and purpose-built to deliver web content only. As a result, while CDNs bring a lot of locations to the table, they do not bring the necessary levels of compute for security processing, only providing limited services like distributed denial of service (DDoS) mitigation and basic web application firewalls (WAFs).

Other limitations include a lack of capacity management tools for customers (as with the public network cloud option) and a significant network design problem. Whereas CDNs seek to absorb requests (for content distribution) or denials (for DDoS mitigation) at the edge to offload traffic from the primary server, a security cloud’s policies will only block a very small portion of requests and responses. To a first approximation, an SSE platform delivers all of the traffic that it receives, absorbing none. That’s the very definition of failure for a CDN, whether being used for content distribution or DDoS mitigation.


Finally, be cautious of CDN providers offering in-house SSE platforms. Take time to fully appraise the security capabilities of each solution for both functionality and efficacy.

Purpose-built private cloud

A final architecture choice is to bring together the diverse network presence of a CDN with the elasticity of a public cloud and a focus on high-performance security services. A purpose-built private security cloud delivers the SASE architecture necessary to complete the solution.

This approach delivers against the “a lot of compute in a lot of places” requirement for an SSE platform, with much of this processing taking place at the network edge (i.e., close to users). The purpose-built, private cloud model is by far the best suited to delivering on an access and security cloud.

While utilizing a purpose-built private cloud does add complexity, including the need to manage interactions among multiple networks, implement an elastic service, and manage complex and sophisticated operations, the benefits of a private cloud model show that the effort is more than worthwhile, delivering:



- Complete control over routing, peering, and traffic engineering.
- Security processing close to end users.
- Scalability to meet variations in demand.
- More distributed computing.
- More effective task-focused peering and routing.

The Netskope NewEdge platform is

a purpose-built private cloud



A common customer experience with Netskope NewEdge is that the **full “SSE” configuration** typically exceeds the end-to-end performance of a **“no SSE” configuration**.

Delivered as a service in a cloud consumption model, Netskope Intelligent SSE is an **immediate outcome-based solution** that meets organizations’ immediate business needs around security without performance trade-offs, and provides future-proof architecture for scalability to address more users, additional services, and expanded coverage.

First and last mile latency

Utilizing the public internet and public clouds typically comes with a trade-off, with a need to choose between performance, availability, and security. Especially for networking, infrastructure, and operations leaders, latency is a well-understood measurement of network performance, and any SSE platform must drive toward the lowest possible latency in every step of a network packet's journey.

To mitigate this trade-off, an extensive peering and interconnection strategy is needed in conjunction with extensive point of presence coverage previously discussed. This should provide for the fast on-ramping of traffic to the security cloud at the "last mile" for remote users and branch sites alike. The strategy should also enable single-digit millisecond latency for the vast majority of knowledge workers globally.

Similarly, to optimize latency from any SSE platform to websites, SaaS applications, or workloads in the public cloud (the "first mile"), architects should look for providers offering single-digit milliseconds of latency to Google, Microsoft, and other leading public clouds, including AWS, Azure, and GCP from every region.



Netskope NewEdge achieves low latency through being carrier neutral and maintaining extensive direct peering relationships as well as interconnects with leading CDNs and internet exchanges.

Today, Netskope NewEdge has

3,000+ network
adjacencies

to **600+** unique Autonomous
System Numbers (ASNs).



These figures were validated in a recent test performed using the third-party Catchpoint service, the results showing that every NewEdge data center is typically within **single-digit milliseconds** of Google, Microsoft, and leading public cloud providers, including AWS, Azure, and GCP.

Processing latency

The last aspect to consider when looking at overall latency is the speed of security traffic processing and services like next-generation secure web gateways (SWG), cloud access security brokers (CASB), or Firewall-as-a-Service (FWaaS). Architecturally, factors like single-pass architecture and containerized microservices, as well as custom data center racks relying on bare metal servers and the highest-performing networking equipment, can play an important role in making traffic processing as efficient as possible.



In 2021, Netskope announced an **industry first** set of service level agreements (SLAs) addressing non-decrypted and decrypted transactions for traffic processing latency inside the NewEdge data centers.

Netskope commits to a

less than **10 millisecond SLA** for non-decrypted traffic
and **50 millisecond SLA** for decrypted traffic.



When comparing SLAs across vendors, architects should consider the small print, specifically whether the SLA applies to the “monthly average” or “monthly 95th percentile.” The latter is a much more stringent SLA and the one used by Netskope for NewEdge customers.

WAN interconnect

Although organizations may choose to connect directly from the user client to an SSE platform, it is common to extend a WAN to directly interconnect with the SSE platform, steering traffic directly from within the WAN.

If you are considering a dual-platform approach, where you maintain existing SD-WAN infrastructure, look for SSE platforms that offer easy integration and tunnel management via GRE or IPSec tunneling protocols. SSE platforms should offer automation options via APIs and support for established primary and secondary tunnels for auto failover on a per-site basis. Any platform should also provide a full overview of configured tunnels and real-time status of tunnels and throughput, to allow easy monitoring of all interconnects.

While Netskope fully supports third-party SD-WAN solutions, for organizations looking for a single-vendor SASE solution, **Netskope's Borderless SD-WAN solution** is directly integrated with the SSE platform, offering one-click tunnel creation and integrated monitoring of the full platform.



Netskope also offers SD-WAN Software Client functionality through the single endpoint client, offering steering for web, both cloud and local apps, along with any SD-WAN network location, offering a truly integrated connectivity and security solution for remote workers.

Organizations deploying Netskope's Intelligent SSE platform typically find that by routing public cloud traffic through Netskope's SSE platform via SD-WAN tunnels and employing QoS can retire private connections to public cloud providers, including Microsoft ExpressRoute, AWS Direct Connect, and Google Cloud Interconnect links. Only made possible by the first and last mile performance of Netskope NewEdge.



Resiliency

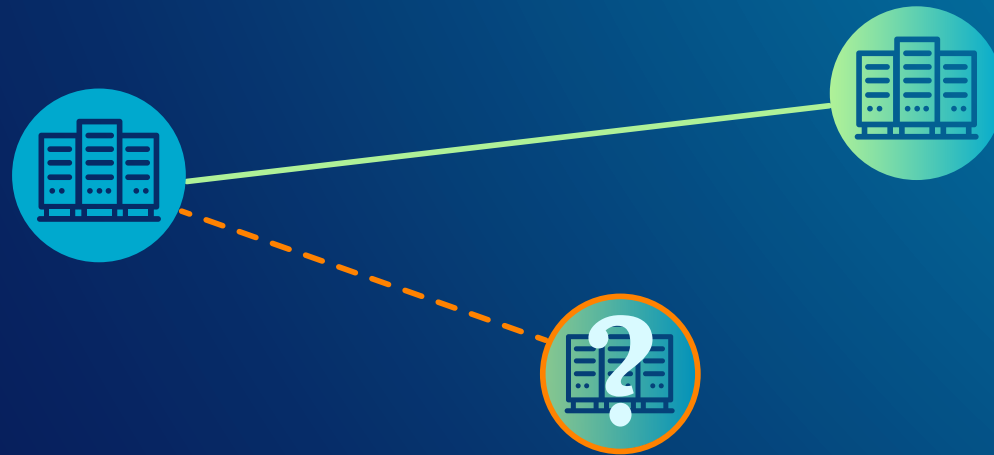
Some infrastructure and operations teams fear that outsourcing network and application access to a third-party platform risks a loss of control, poor service availability, and limited resilience.

To address these concerns, many service providers are introducing availability SLAs. For example, with the Netskope NewEdge offering, customers benefit from a five nines (99.999%) SLA related to availability of inline services.* However, it's imperative that architects understand whether the vendor can actually back up its SLA claims. Architects should ask a number of key questions, including:

- How does the service provider handle outages at the data center level, whether through planned maintenance or unplanned incidents?
- What happens during major environmental events, like earthquakes, floods, or fires?
- What smart and automated failover capabilities do they utilize and how is this tied into global service monitoring to proactively address unforeseen events?

* NewEdge POPs in mainland China offer alternative SLAs. For more information on the specific services included, please refer to the "Netskope Support and Service Level Terms" document.

Netskope's inline capabilities are configured with NewEdge Traffic Management so that anytime a data center becomes unavailable for any reason user traffic is **automatically directed** to the next closest available data center to maintain service levels.



Netskope's philosophy is to overprovision data centers so that performance and service availability is maintained following a data center or regional outage. Additionally, our service monitoring **detects potential issues early**, so customers are not impacted and robust security protections remain in place at all times.

Summary

An effective access and security cloud requires more than the ability to apply computing to network traffic. To deliver both performance and security, an access and security cloud needs substantial computing near users, excellent visibility into capacity, well-managed peering, and an architecture that supports the low-overhead transit of traffic. In short, an access and security cloud worthy of the name requires careful network design.

Designing and implementing a security cloud is not the work of a moment. It requires significant investment and a team of people with the experience needed to execute and operate an extensive custom-built platform. This is why Netskope has invested over \$150M to date in building the NewEdge platform, designed by a team of people who were also involved in building the world's largest public clouds and content delivery networks.

With Netskope NewEdge, we offer a purpose-built global security private cloud that addresses these requirements more effectively than any alternative technical model. Not only that, Netskope works closely with the security teams of our customers to build a unified architecture that meets their unique needs.



To learn more about NewEdge and how you can progress a new access and security architecture, visit <https://www.netskope.com/platform/newedge>.

About Netskope

Netskope is a leader in Secure Access Service Edge, redefining cloud, data, and network security and helping organizations apply zero trust principles. The Netskope Intelligent Security Service Edge (SSE) platform is fast, easy to use, and protects people, devices, and data no matter where they are. Netskope helps organizations reduce risk, increase effectiveness, and gain unparalleled visibility into all cloud, web, and personal application activity.

Thousands of customers, including more than 25 of the Fortune 100, trust Netskope and its powerful NewEdge network to mitigate threats and address technological, organizational, network, and regulatory changes.

