

Your Network of Tomorrow

Four Principles for Modern Network Design



Flexible

Secure



Table of Contents

THE SECULAR FORCES DRIVING CHANGE	5
Cloud Adoption	
Data Expansion	
Hybrid Work	
THE THREE ERAS OF ENTERPRISE NETWORK DESIGN	6
The First Era: Physical Networks (3-Tier / Spine-Leaf)	7
The Second Era: Virtualized Networking	8
The Third Era: Cloud-Delivered Network and Security Services	9
SASE—A Practical Implementation of Cloud-Delivered Network and Security Services	10
FOUR PRINCIPLES FOR IMPLEMENTING CLOUD-DELIVERED NETWORK AND SECURITY SERVICES	11
PRINCIPLE 1: ELIMINATE SOURCES OF COMPLEXITY	12
Simplify Your Architecture	
Modernize Outdated Designs	
PRINCIPLE 2: GO DIRECT TO INTERNET	13
Distribute Access to Security Services to Reduce Latency	15
SASE Service Delivery	16
PRINCIPLE 3: REPLACE CLOUD INTERCONNECTS	16
Using SASE Peering as an Alternative to Cloud Interconnects	18

Table of Contents

PRINCIPLE 4: REMOVE IMPLICIT TRUST	19
Evolve the LAN	20
Model the LAN as the Guest Network for All Users	21
Intermingled Managed and Unmanaged Applications	22
Using SASE as the Perimeter for Managed Applications	
Example Mapping: Managed vs. Unmanaged Applications	23
Broker Access and Enforce Authentication / Identification with SASE Rather Than Relying on NAC/802.1x/VPN	24
Using SASE to Enforce Identity-Based Policy	25
Rethink the DMZ	26
Architectural Alternatives to DMZ	27
ABOUT THE NETSKOPE SASE PLATFORM	28
Netskope SASE	28
Using Netskope for Your Network of Tomorrow	29

Foreword: Note to the Audience

We are seeing many customers approaching a fork in the road in their networking goals, and asking for help on how to adapt today's architecture for tomorrow's needs. We've worked with thousands of customers on this journey, and this guide summarizes the four core principles for networking teams to build your network of tomorrow.

Information is the lifeblood of business. In order to maintain competitive advantage, organizations design enterprise networks to deliver fast and reliable access to information from the applications in which it lives.

Today we are in an era of rapid change, where both internal and external factors are driving conditions where flexibility, rather than efficiency alone, plays a crucial role between success and failure. This is a dramatic change from the days of optimizing for cost-effectiveness to improve the bottom line. To thrive under such conditions, management must have IT strategies that are cost-effective, flexible, and rapidly adapt to new imperatives at the speed of business.

We have seen how inflexibility manifests in recent years. At the onset of hybrid work, many organizations were not able to adapt quickly, with capital tied up in network designs that presumed users and applications were primarily on the LAN. Under a dramatically new set of requirements, the organization scrambled to adapt.

The pressure is on to make changes to the network, but what are the key priorities? Use this guide to plan your path toward a faster, more secure, and more resilient network designed for the applications and users that you support.

THE SECULAR FORCES DRIVING CHANGE

In order to determine how the network should change, it's important to understand the bigger picture on why it must change at all. The following secular forces are reshaping the landscape for the way we live and work, and directly impact the IT team's responsibilities. Each factor is creating pressure to redesign the network, and not in small, incremental ways. Taken together, it's driving an imperative for transformation.

Digital transformation has firmly taken hold in the majority of organizations around the world. A recent survey of CIOs shows that 60% of companies will continue to make significant investments in digitalization to enhance competitive capabilities, enable greater business agility, and aid in decision-making. ¹

Cloud Adoption

One aspect of this evolution is that enterprise applications and data are increasingly moving out of corporate networks and data centers and into the cloud. According to Gartner, 70% of all enterprise workloads, up from 40% in 2020, will be deployed in cloud infrastructure and platform services by 2023. ² What's more, over 80% of all enterprise traffic is destined for the internet and 53% of all web traffic is cloud related. ³

Over 80% of all enterprise traffic is destined for the internet and 53% of all web traffic is cloud related.

Data Expansion


Another critical factor is the rapidly expanding volume of data being generated. Between 2020 and 2025, the amount of data in the world will increase from 57 zettabytes (ZB) to an astounding 175 ZB. ⁴ More data is being collected and shared across more points of access than ever before, and this information is widely distributed across both networks and clouds, as well as a number of managed and unmanaged devices. Without purpose-built protection, these vast quantities of distributed data are quite vulnerable— more than one-third (40%) of organizations experienced a cloud-based data breach last year. ⁵

Hybrid Work

A third factor contributing to turbulence is that a significant proportion of the user population will continue to work outside of a traditional corporate office location with the expectation of being able to access information from any device and location, all without having to make trade-offs between security. 50% of the U.S. workforce is expected to continue to work from home long term. ⁶

With these secular forces driving the reasons for change, it's now becoming clear what we must solve for. Toward this end, the very goal of the network and how it should be implemented is driving the third era of enterprise network design. It's now clear that the existing network is optimized for an entirely different set of conditions, and we must rapidly solve today's challenges. In fact, adherence to old network and security architectures can hinder the ability to capitalize upon the benefits of the cloud and hybrid work, and it can also miss the importance of getting data protection across the entire enterprise landscape right.

THE THREE ERAS OF ENTERPRISE NETWORK DESIGN



There have been three major eras that radically reshaped thinking in network designs:

1. Physical Networks (3-Tier Hierarchy / Spine-Leaf Switching)
2. Virtualized Networking
3. Cloud-Delivered Network and Security Services

The baseline architecture that's at the heart of most enterprise-level networks is the 3-tier network and its variants based on spine-leaf switching. As a time-tested, well-understood model for network service delivery, it remains relevant even today.

We are currently well into the mainstream adoption of the second era of network design, virtualized networking, which addresses the growing need for networking and security services insertion for east-west traffic. Instead of sending east-west traffic to the core, virtualized networking implements switching, routing, and security services in the hypervisor or a bridge mode virtualized firewall; virtualized networking thus offloading (but not eliminating or replacing) the load on core network services.


One of the notable aspects of the first two eras of network design is that one does not replace the other. The implementation of virtualized networking coexists with the physical networks that preceded it, but it does not reduce or eliminate any of its design. Therefore, in terms of functionality, virtualized networking provides a tremendous benefit, but remains on top of the underlying physical network.

With the advent of cloud-delivered network and security services, the third era of modern network design is now upon us. Unlike the first two eras, Cloud-delivered network and security services provide organizations with a tremendous opportunity to take complexity out of enterprise networks. Instead of the organization bearing the entire capital and operational costs of running enterprise network services, the third era uses cloud-delivered services to augment the WAN. In the process of adopting cloud-based services, there are many opportunities to make the enterprise network more efficient, reliable, and secure by subtracting, rather than adding, complexity. This is possible because the networking and security services can be implemented as cloud delivered services without having to reinvent or re-engineer the WAN design.

In this guide, we will explore the arc of network design and provide guidance as organizations move forward on their network journey.

The First Era: Physical Networks (3-Tier / Spine-Leaf)

The basic concept of the 3-tier network concentrates high-powered networking services in the core network, while using layers of distribution to establish connectivity everywhere the organization operates. From the perspective of networking, it's a proven architecture for connectivity.



Example of a Physical Network

The traditional 3-layer design attempts to minimize costs by maximizing utilization of centralized, powerful (and expensive) routers and firewalls in the core. As such, the core layer connects the WAN segments and serves as the demarcation between trusted and untrusted networks.

Security implementations into the 3-tier network introduces a number of complexities. Security must inspect without slowing down the network. It introduces architectural challenges such as where to insert into the physical architecture. As requirements evolve, additional products for firewall, intrusion prevention, and data loss prevention services are inserted into the network design, creating more operational complexity and points of failure.

With the evolution of the data center and the ever-expanding need for access to applications, the physical network in enterprises shifted to the spine-leaf architecture. Spine-leaf architectures reduce latency in east-west traffic by interconnecting leaf switches in a full mesh. In practice, spine-leaf simplifies one aspect of the architecture (by collapsing three tiers down to two) but remains as complex or even more complicated when thinking about security. Even with improvements to switching, where should inspection be inserted?

The first era of enterprise networking excels at delivering connectivity, but security still remains difficult to insert. It was designed to facilitate rather than limit access. Isolating hosts is easy, but connecting isolated hosts and implementing precise controls and content inspection is difficult. The need for better isolation sets the tone for the second era of enterprise networking.

The first era of enterprise networking excels at delivering connectivity, but security still remains difficult to insert. It was designed to facilitate rather than limit access.

The Second Era: Virtualized Networking

Several important developments helped to catalyze and define the era of virtualized networking. First, virtualization radically changed the structure of data centers, increasing the host density while also introducing machine-to-machine traffic that never hits the physical wire. Second, the risk of malicious insiders and attackers piggybacking on compromised machines brought more attention to the problem of lateral movement across systems at similar trust levels. Third, zero trust design principles drove requirements for more isolation and granular access controls than the flat network or traditional network segmentation.

Virtualized networking was the birth of the second era of network design. Rather than attempting to force all traffic through the core, organizations began to use a number of overlay technologies to add networking controls such as segmentation without relying on changes to the underlying network. In addition, by virtue of using the overlay, the virtualized network improved upon the physical network by making it possible to insert security inspection in east-west traffic without forcing traffic on the wire. The implementation of these services could be done by using virtualized networking to route traffic through virtualized firewalls or adding security services that hook into the hypervisor.

In many ways, however, virtualized networking improved the organization's ability to perform more types of security, but it still faced a complexity problem. Virtualized networking is an overlay, and the complexity of both the physical and virtualized networking services coexist.

The Third Era: Cloud-Delivered Network and Security Services

The first two eras of network design were entirely endogenous, meaning that the IT department designed, built, and operated the infrastructure that provided the networking services. However, what happens when the resources you connect are exogenous to the enterprise network? This is precisely what happened as the effects of cloud and hybrid work changed the very nature of what the network needed to connect and secure.

These conditions introduce a compelling question: “Could key networking and security services be delivered from the cloud, and operate in conjunction with the classic enterprise network?”

Many, if not most, enterprise applications involve the cloud in some manner. Forcing traffic through the enterprise network to reach cloud applications is problematic, for there are a limited number of egress points that the organization operates, and a fixed amount of firewall capacity. In most cases, forcing traffic through the enterprise security stack creates suboptimal routing that increases latency and damages the user experience as frustrated users wait for their applications to respond.

By using networking and security services from the cloud, organizations can deliver access to applications to users wherever they are located without forcing traffic to on-premises resources. However, instead of treating the access to cloud as a separate use case, what’s not apparent is that the third era provides an opportunity to greatly reduce the complexity inherited from the first two eras of network design.

Treating the cloud as an extension of the enterprise network (i.e., a peer networking+security delivery network) creates an opportunity to eliminate sources of complexity that were baked into the on-premises network. In fact, cloud-delivered network and security services decouples the need to “add” more stuff to your network, because it eliminates the need to force traffic through an appliance to implement new services.

The cloud-delivered model reduces the networking problem down to a first hop to the cloud provider’s data center, which delivers security microservices that can be enabled without further network changes. By decoupling the delivery of security from the network, the organization can run leaner, more reliable enterprise networks with security delivered from the cloud.

“Could key networking and security services be delivered from the cloud, and operate in conjunction with the classic enterprise network?”

SASE—A PRACTICAL IMPLEMENTATION OF CLOUD-DELIVERED NETWORK AND SECURITY SERVICES




Secure Access Service Edge (SASE), pronounced “sassy,” is a cloud-based architecture that delivers network and security services meant to protect users, applications, and data. Given that many users and applications no longer live and operate on a corporate network, access and security measures can’t depend on conventional hardware appliances in the corporate data center.

SASE delivers the necessary networking and security capabilities in the form of cloud-delivered services. Instead of routing traffic to an appliance for security, users connect to the SASE cloud service to safely access and use web services, applications, and data with the consistent enforcement of security policy.

FOUR PRINCIPLES FOR IMPLEMENTING CLOUD-DELIVERED NETWORK AND SECURITY SERVICES

As we enter the start of the third era of enterprise networking, how should network architects think about what their future state should look like?

This guide will help you along the way. These four principles help organizations understand where to implement meaningful changes to improve their networking services:



PRINCIPLE 1: ELIMINATE SOURCES OF COMPLEXITY

SASE changes the game for network design and it opens the doors to new possibilities.

Simplify Your Architecture

A well-designed network should always have extra capacity because operating near its breaking point is a recipe for failure. A resilient network has to be able to absorb shocks to utilization, such as when most of the company's salesforce is at headquarters for a sales conference, or when the company migrates to or from a remote work policy. Even under the normal follow-the-sun business operations cycle, any given location on the network might be operating close to idle for 16 hours outside of the standard workday. But extra capacity comes at a cost because underutilized networks still require equipment, services, and staffing that are ideally unused.

In a similar vein, reliability comes at a cost as well. Redundancy doubles the expenses for network equipment as well as failover services, which are ideally NOT being used if they are cold spares or passive failovers.

As a consequence, every service inserted into the network requires overprovisioning and overengineering to account for stability, usability, reliability, and scalability. This is not a bad thing, for organizations should make sure that there is capacity and reliability in the design. The question, however, is whether insertion into the network is the best way to implement such services. If one could "shift" capacity and the related overprovisioning and engineering to the SASE layer, then you would only pay for what you use when you actually use it.

A well-designed network should always have extra capacity because operating near its breaking point is a recipe for failure.

Simplification increases the availability and resilience of the existing network. It is easier to achieve carrier-grade, dial-tone reliability and elasticity if there are fewer things to break.

To simplify the network, organizations should think about how to streamline the network design, get services and capabilities rolled out faster, address new offices, branches, or remote workers, go direct-to-net for cheaper connectivity and faster access to SaaS, or leverage new technology like SD-WAN. From this point of view, the enterprise network should focus on simple, fast, and reliable networking, while leveraging the SASE layer for the secure connectivity to the cloud.

Modernize Outdated Designs

Within the first two eras of enterprise networking models, inline security services required installing inline physical and virtual appliances. Over time, the accumulation of different security services made routing inefficient, such as the complexity created by proxy chaining legacy security products or the need to tunnel traffic to reach a distant security stack. These practices often remain in use today, with many vendors developing products that use inefficient chaining and forwarding behind the scenes.


To modernize the architecture, use SASE as a peer network to implement services rather than constantly rewiring the network. For example, instead of implementing multiple gateways strung together to address the multiple lanes of internet, web, and cloud traffic, the SASE layer delivers SWG, Cloud Firewall and CASB that can be implemented from the same first hop. Deploying the first inline product makes it possible to add others by turning on services rather than inserting more security appliances into the network.

Note that as security service delivery shifts to the SASE layer, the underlying network becomes more secure by virtue of eliminating routes and policy exceptions. That's because legacy network access controls relied on implementing policy decisions in the network, whereas the modern network eliminates them. For example, legacy network access controls typically rely on firewall policies to network segments or zones. Every "allow" can have unintended consequences, especially when considering the coarse grain of network segmentations. By taking out the policy decisions from the network layer, the attack surface shrinks and collapses to the fine-grained policy decisions in SASE.

As organizations evolve to add new security capabilities, such as data protection and threat protection, such services can be activated from within the SASE platform.

PRINCIPLE 2: GO DIRECT TO INTERNET

Wide area networks use hub and spoke designs for several reasons. In part, it used to be necessary to route all traffic through the core to reach the internal data center. It also was a logical design to insert security at the limited number of egress points.




With applications moving to the cloud, enterprises started efforts to adopt a direct to internet approach. Using an internet connection at the branch, traffic bound for the cloud and internet could be offloaded from the MPLS network. The problem is that low-cost branch firewalls rarely have the same security as the perimeter firewall, thus making security policies inconsistently enforced. Even if the organization buys the same perimeter firewall for the branch, the maintenance work on the branch boxes strains overburdened network teams. Or worse, it's put in the hands of non-technical branch employees with no experience managing firewalls or controlling the physical security of the devices. As a consequence, solving one problem in the network creates another, possibly worse, operational problem.

Using SASE for Direct to Internet

Several aspects of the SASE model make it possible to achieve simplicity, modernization, and optimization in branch networking by implementing direct to internet for branch locations. Note that the two main hurdles toward replacing MPLS are 1) instability and reliability issues of the open internet and 2) enforcing consistent security without expensive, high-maintenance gear at the branch. These challenges are overcome through the delivery of networking and security through SASE.

To overcome the instability and reliability issues of the open internet, use multiple connections, including low-cost internet, orchestrated and managed through SD-WAN. SD-WAN maintains session stability even when congestion or link failure occurs, as it can ensure that the session stays alive as it makes the respective adjustments to optimize performance. As this happens transparently, the SD-WAN connections keep applications running without disruption as traffic takes advantage of direct to internet.




To overcome the problem of inconsistent security, organizations use the implementation of SD-WAN with SASE to deliver security rather than relying on the on-premises firewall. Thus, by offloading security services from the branch, the primary job of the local device is much simpler. The local SD-WAN device establishes the networking to first hop into the SASE cloud to establish visibility and security as organizations use applications. Thus, organizations do not need to maintain expensive firewalls at the branch.

Note that the application of SD-WAN technologies is not solely limited to the branch. The growing dependence upon cloud applications, especially ones with real-time communications such as collaboration and video functionality, also have ramifications on the connectivity that users have when working remotely. As your organization evaluates SD-WAN, consider whether it makes sense to use multiple links from the endpoint to the internet (such as broadband in conjunction with a 5G modem) to improve the stability of critical applications.

Distribute Access to Security Services to Reduce Latency

By and large, most organizations have physical and financial limits to the number of data centers they can operate and the number of egress points they can support. For example, a large organization might implement 4 major data centers to cover the world (2 in North America, 1 in Europe, and 1 in Asia), and require all branch locations to connect to one of these hubs. It's cost prohibitive to add the 5th or 6th site, thus leading to a coverage ceiling where the company simply tolerates suboptimal routing until it can cost justify opening another data center to relieve the burden.

The consequence to the coverage ceiling is that only a portion of the organization's population is near an egress point, and everyone else accepts some degree of latency. The further the user or office is from the data center, the more latency users experience, which is especially problematic with highly mobile workforces.



Most current count available at <https://trust.netskope.com/>

SASE Service Delivery

With SASE, the burden of building out and provisioning coverage shifts to the SASE vendor which delivers the coverage that's provided to its customers. This makes it possible to break the coverage ceiling because the SASE provider typically builds out in far more regions than the average enterprise. The net result is that the organization has more global coverage than it could possibly support on its own, without the burden of managing global data centers.

PRINCIPLE 3: REPLACE CLOUD INTERCONNECTS


The networking services that organizations use to connect to the internet, build their WAN, and extend access to the cloud rely on services provided by local providers. However, these services are not universally consistent around the world. The service providers will vary from market to market, the level of service will differ, and price points create both technical and economical limits to the organization's ability to support different geographies.

The link to the internet is usually the most cost effective, with the caveat that internet services are not as reliable or implicitly private as a dedicated link. With the explosion of cloud computing, companies had a growing number of business-critical applications that were external to their network. Given the uncertainty that frequently came from using the internet to access the cloud, cloud interconnects such as ExpressRoute for Azure, Direct Connect for AWS, or Express Connect for Salesforce became popular to support critical cloud applications.

The original thinking was to link the cloud->data center communications in a similar way to data center->data center communications.

The first wave of cloud interconnects established the model of hybrid cloud architecture (i.e., using public cloud as an extension of the on-premises private cloud data center). The original thinking was to link the cloud->data center communications in a similar way to data center->data center communications.

Over time, the use of interconnects expanded. For example, Salesforce is every bit as mission critical for organizations to conduct transactions with their customers as an on-premises application, and thus requiring a dedicated link for speed and reliable access from the WAN. The problem is that organizations have many SaaS apps, many of which are business critical, so which ones should have an interconnect and which ones route over the open internet? In addition, global coverage is also an issue. If a regional office does not have a local option for cloud interconnect, should the organization internally route the traffic to a location that does, or should it go ahead and use the open internet?



Traditional IT-Managed Cloud Interconnect Over HA Dedicated Links

The security argument is sometimes used as a justification to add more cloud interconnect links. For example, if all traffic to an application is routed over a dedicated link, then internet access could be eliminated, thus reducing exposure to pre-auth exploits or attempting to credential-stuff the authentication. If the origin of the traffic comes from the internal WAN rather than the internet, then access to the application would be at least as secure as the WAN.

This security is only relative, however, because it still remains possible that a bad actor is operating on the WAN and can laterally move to the interconnect. As such, organizations continued to make the cloud interconnect architecture even more complicated by adding bastion hosts in front of the cloud interconnect and requiring external users to use VPN in order to reach it.

What's become clear is that connecting the WAN to the cloud over interconnects is not scalable or architecturally sound. A better approach is to leverage the security and networking from the SASE layer for highly performant and secure access to the cloud applications in different parts of the world.


With the explosion of cloud computing, companies had a growing number of business-critical applications that were external to their network.

Using SASE Peering as an Alternative to Cloud Interconnects

Instead of sourcing and maintaining cloud interconnects from a multitude of regional service providers, consider using SASE peering with cloud providers for reliable, secure, low-latency access to cloud applications. This approach greatly reduces cost/complexity by offloading the complicated service delivery architecture to the SASE provider.

Not all SASE vendors are equal in this regard. For example, one vendor might optimize purely for the processing time within its own data center, without thinking about how to deliver the best round-trip time to the user. As a result, not all exits from the SASE data center are equal. It's important to validate that your SASE vendor of choice optimizes security and networking performance.

When done properly, IT steers traffic to a SASE vendor's data center to leverage the vendor's interconnects and peering to the major cloud providers as well as the rest of the internet. It can deliver consistent coverage and speeds compared to the variances that IT normally faces with regional network providers.



PRINCIPLE 4: REMOVE IMPLICIT TRUST

Perimeter security separates the external, untrusted entities from the internet from those on the internal network. This model is sensible at the coarsest level of segmentation (internal versus external) because it can be used to control policy between external actors and internal resources, and internal actors with external resources.

However, once the user is inside the perimeter, there are few controls that gate access to internal resources. As a byproduct, the internal network often has excessive levels of implicit trust that presumes all internal users are “trustworthy.” Of course, trustworthy is a relative term given that it’s presumed, but not guaranteed, that internal users will not harm the company.

To correct the problems of implicit trust, many companies have tried to insert a number of security technologies into the network to assert control:

1. NAC to enforce device-level authentication and the application of L2 policies for network access
2. VPNs to connect remote devices to the local network over a tunnel
3. Network segmentation for coarse-grained separation of network functions and security zones, such as the data center versus LAN, or web server to database server
4. VLANs for logical separation of functional groups, such as separating marketing from accounting
5. Host-checking policies to check device config and patch level
6. Two-factor authentication to reduce the risk posed by stolen credential

All of these technologies, while well intended, aim to remove some implicit trust, but add more complexity to the network as a consequence. Even worse, security gaps remain. For example, if device authentication and user identity policies are separately enforced, then it’s possible for a remotely controlled device to port scan the network and identify servers with unpatched vulnerabilities even without any identity credentials. And many of these security technologies, such as NAC and VPN, are not relevant when accessing cloud applications.

In terms of design goals, eliminating implicit trust actually makes the network simpler to operate, for a number of reasons.

Today, what’s become clear is that you can’t add enough security to eliminate implicit trust. You have to architect the trust out of the network in the first place. In terms of design goals, eliminating implicit trust actually makes the network simpler to operate, for a number of reasons.


- Policy controls can be implemented from the security layer rather than trying to make allow/deny decisions in the network design.
- The networking is more reliable if there are fewer security appliances or services that need to be managed.
- Many outbound paths to different network segments and applications can be consolidated down to a first hop to the security services layer.
- Many inbound paths can be eliminated entirely, thus reducing the exposed attack surface layer and the potential for policy mistakes or abuse of stolen credentials.

Networking teams can take steps to eliminate sources of implicit trust:

- Evolve the LAN—minimize or eliminate external access and do not grant any extra privilege for connecting from the LAN
- Route all managed application traffic through SASE
- Broker access and enforce authentication / identification with SASE rather than relying on NAC/802.1x/VPN
- Rethink the DMZ

Evolve the LAN


The LAN is the front door for local hosts to access the corporate network and data center, and yet it remains a problematic source of implicit trust. For years, the only protection for the LAN was the guard at the front desk and the badge reader to the building. However, once inside the building, a device could connect to the LAN with broad levels of network access, typically in excess of what the user actually needed. While there have been efforts to control devices on the networks, many enterprise networks are extremely vulnerable to attacks from within, either by an insider or an attacker that is operating from a compromised device. Once on the LAN, a bad actor typically faces little resistance mapping out the rest of the network and compromising other hosts.



Are broad levels of network access on the LAN even necessary today? In the past, it was common to use the LAN to reach applications in the data center and enable workgroup functions such as file shares, collaboration tools, and real-time communications. Today, by and large, such services have evolved to cloud-based equivalents and no longer require extensive, excessive LAN services.

Model the LAN as the Guest Network for All Users

Fortunately, the model for removing implicit trust out of the LAN already exists in most enterprise networks. It's the guest network. The guest network's function is to provide internet connectivity to non-employees, such as contractors and guests. It does not permit access to the data center or other resources on the LAN.



Today, internet connectivity is really all that's needed for the majority of on-premises employees to use their applications. With hybrid work, the office is no different than the remote user, given that the applications are accessed in the cloud.

As such, there is no need to grant users more access to the network just because they are working from the office. In fact, with hybrid work, the office networking starts to resemble shared workspaces where users from different organizations are using a shared network with no implicit trust between hosts.


By virtue of enabling employees to work from anywhere, the distinction of being on a trusted network versus untrusted network is disappearing. Working from home, a coffee shop, or a shared workspace, the user spends the majority of their time on untrusted networks. Placing employees on an untrusted LAN or guest network while at the office is no different.

By virtue of enabling employees to work from anywhere, the distinction of being on a trusted network versus untrusted network is disappearing.

Intermingled Managed and Unmanaged Applications

In the past, it made sense to delineate the perimeter as being the border between the trusted and untrusted network. With data, the perimeter now shifts to enforce controls between trusted and untrusted applications, and you want to ensure that your data does not move to places where you cannot control it.

However, conventional security appliances are not well-suited to make the distinction between which application or instance is managed or not. The growing number of exceptions to visibility creates a dangerous condition where managed and unmanaged applications are indistinguishable from one another, and no controls are in place to stop unwanted data movement



Using SASE as the Perimeter for Managed Applications

Instead of trying to force managed traffic through a data center firewall, first think about what your managed applications need for security and where the applications live.

For example, if you classify your applications as managed / unmanaged, then it becomes quite clear that many security services need to be delivered across the Data Center, Public Cloud, and SaaS.


The growing number of exceptions to visibility creates a dangerous condition where managed and unmanaged applications are indistinguishable from one another, and no controls are in place to stop unwanted data movement.



Example Mapping: Managed vs. Unmanaged Applications

APPLICATION TYPE	LOCATION	EXAMPLE APPLICATIONS	KEY SECURITY REQUIREMENTS
MANAGED APPLICATIONS	Data Center	Oracle Database	• Secure Access • Data Protection • Behavioral Analytics
	Public Cloud	AWS / Azure	
	SaaS	Office 365 Google Workspace Salesforce Workday GitHub	
UNMANAGED/ PERSONAL APPLICATIONS	SaaS	Microsoft 365 Google Workspace Dropbox Zippyshare	• Risk Scoring • Risk-Based Policy • Instance Detection • Threat Protection • Data Protection • Behavioral Analytics • End-User Coaching
	Internet	Web Non-Web (FTP, SSH, RSH, etc.)	

The path forward is to think about how to control the data in your managed applications and deliver the necessary security wherever needed. By routing all of your managed applications through SASE, you can thus ensure there are secure access, data protection, and behavioral analytics capabilities. In a similar vein, all remaining application traffic can be treated as unmanaged or personal, with policies in place to prevent unwanted data movement.



Broker Access and Enforce Authentication / Identification with SASE Rather Than Relying on NAC/802.1x/VPN

The marriage between identity and networking has a long and rocky history. Today, networks implement identity policies in different layers of the network connection. Metaphorically, each security measure acts as a gate before allowing a connection to succeed, with little to no knowledge of the policy decision before or after it. For example, NAC establishes whether a device is permitted to connect to the access layer of the network, but at layer 2 it would not be able to determine whether the user on the device is authorized to access a given application. Thus, the NAC's field of view is solely to determine if a device is allowed on network, but not what it can do.



In effect, these technologies are attempting to mitigate implicit trust in the network by creating additional barriers to protected resources. It remains problematic for several reasons:

1. Any level of access leaves the network open to abuse. For example, just being connected to the network without application credentials leaves the door open to surveillance, port scanning other hosts, checking for unpatched software, stealing credentials, and pre-authentication exploits on vulnerable systems.
2. Placing identity controls in the network layers is not useful when neither user nor application are on the network. It would be inefficient to route traffic through the network to get these identity services. Therefore, the delivery of identity controls to managed applications must account for work-from-anywhere scenarios.


Using SASE to Enforce Identity-Based Policy

Instead of gating network connections to an application, it makes more sense to apply zero trust network access (ZTNA) to broker connections with identity criteria to deliver user->application controls and build connectivity through the SASE infrastructure.

As mentioned earlier, by reducing the trust placed in the internal LAN, the organization dramatically reduces the attack surface. With all users limited to internet-only, and no internal routing to managed applications, then there is no implicit trust to claw back.

The bridge between users and applications is the SASE overlay to implement ZTNA, which not only checks the identity and criteria policies to allow access, but it also connects two outbound connections. The beauty of the zero trust network access approach is that your network firewall policies become much easier at the data center as well. On the application side, simply block all inbound traffic. This means that there's no gateway to probe, no visible server exposed, and nothing to port scan. If there's no inbound traffic at all, the attack surface disappears.


The beauty of the zero trust network access approach is that your network firewall policies become much easier at the data center as well.



Rethink the DMZ

The network DMZ provides a method to expose internal resources to the public, but is it truly necessary today? At one time, it was a useful and yet imperfect way to expose a managed network to the internet in order to support a number of use cases, including:

- **For public-facing custom applications:** Web servers and application servers placed in the DMZ could be used by the public.
- **For enterprise apps used by employees and contractors:** Applications either running in the DMZ or providing limited access to the data center.



But the DMZ is a dangerous place, as it allows the public to interact with servers that have access to internal resources. Any security vulnerability or policy misconfiguration that's public facing could be the catalyst for a breach. As a consequence, organizations have to commit enormous capital to configure and maintain the DMZ:

- One or more network firewall HA pairs to establish the DMZ network
- Load balancers to support and distribute legitimate traffic
- DDoS filtering in the network as well as upstream to filter out attempts to overwhelm the interfaces of servers in the DMZ
- Intrusion prevention to drop attempts to exploit an unpatched vulnerability
- App firewalls to filter out malicious input to applications such as SQL injection and cross-site scripting

Even a well-managed DMZ with careful monitoring and best practices configuration is subject to exploitation. Any system in the DMZ is a part of the attack surface, which could be externally exploited by unauthenticated users.

Architectural Alternatives to DMZ

Instead of managing and securing the DMZ, it's pragmatic to evaluate whether your applications need to operate from the DMZ in the first place. By and large, the cloud changed the landscape for traditional DMZ applications, especially for public-facing applications. This is because purpose-built services such as hosted applications or SaaS deliver similar or the same application capabilities without having to manage the operating environment.

The question, however, remains about what to do about the enterprise apps for employees and contractors. Custom applications can be moved to the public or private cloud, but access to the application still normally requires some type of publicly exposed method to connect to the application, such as a VPN, bastion host, proxy, or terminal server. While each of these technologies intends to provide secure access to the internal server, there are exposed gateways that remain vulnerable to being port scanned, misconfigured, or exploited.

Using the same principles as discussed in using zero trust network access and SASE to broker connections, organizations can extend access to private applications without requiring a DMZ. Unlike the traditional VPN, which connects the end-user's device into the trusted network, zero trust network access delivers secure access specifically between users and applications.

In an ideal scenario, maximize cost and security benefits by eliminating all ingress traffic. This may or may not be achievable when weighed against the requirements of your organization, but it's clear that even a reduction in the number of systems placed in the DMZ will reduce the attack surface and make your network easier to manage.

By and large, the cloud changed the landscape for traditional DMZ applications.

ABOUT THE NETSKOPE ONE PLATFORM

Netskope One is a converged security and network as a service platform. Through its patented Zero Trust Engine, AI innovations, and the largest private security cloud we make it easy for our customers to defend their businesses and data while delivering a phenomenal end user experience.

The Power of One

Reduce cost and complexity with the only converged platform to deliver one engine, one client, one gateway, and one private security cloud.

AI-Powered Visibility & Protection

Decrypt and decode all cloud, SaaS, and web user activity with patented AI innovations and the incomparable visibility and precision of the Zero Trust Engine to stop data loss and threats.

A Phenomenal User Experience

Accelerate your hybrid enterprise with the industry's fastest and most reliable private security cloud infrastructure and proactively manage the experience from user-to-app and everything in between.

COMPONENTS OF THE NETSKOPE ONE PLATFORM

Zero Trust Engine

The Zero Trust Engine lies at the heart of the Netskope One platform, facilitating zero trust implementations by continuously gathering risk telemetry about users, devices, applications, transactions, and data, to extend precise & comprehensive adaptive trust that minimizes risk while delivering a fast and empowering user experience.

NewEdge Network


Netskope NewEdge is the largest private security cloud infrastructure, featuring ultra-fast data centers with coverage in every country for users and offices to connect to the Zero Trust Engine, as well as extensive peering relationships with the world's leading web and cloud providers to ensure a best-in-class end-to-end user experience.

Netskope One Client

Netskope One Client is the industry's first SASE client, unifying remote user access to web, cloud, and private apps, alongside endpoint DLP and endpoint SD-WAN, resulting in a single agent footprint for SASE and zero trust use cases that greatly simplifies desktop administration and the user experience.

Netskope One Gateway

Netskope One Gateway facilitates branch transformation to a thin branch with the secure, optimized connectivity of Borderless SD-WAN, and a seamless on-ramp to Intelligent SSE and consolidation of multiple disjointed branch appliances to one appliance.



Using Netskope for Your Network of Tomorrow

Netskope SASE helps organizations achieve their network of tomorrow by implementing the principles described in this guide. It implements the security and networking services to provide comprehensive coverage from anywhere your business operates. The Netskope NewEdge Network and Private Cloud is the world's largest, highest-performing security private cloud and powers the Netskope One platform. Architected by industry experts who previously launched and scaled some of the world's largest cloud services and content delivery networks, NewEdge is a hyperscale, SASE-ready architecture that provides customers with unparalleled service coverage, performance and resilience.

Provisioning full compute and all the security services available in every one of our 75+ regions and 100+ centers (as of August 2024), NewEdge ensures zero performance trade-offs. This allows security deployment wherever your people and devices are, eliminating reliance on virtual points of presence or backhauling, resulting in faster connectivity to any application and delivering exceptional user experience.

Our cloud-delivered service model processes traffic with lightning speed and availability. Our 5-9s SLA guarantees the uptime and availability of our inline services. Netskope massively overprovisions the NewEdge network, which is capable of achieving 2 terabits per second at each data center or more than 100 terabits globally. All of this is part of the NewEdge design, which helps our customers simplify their own networks.

The SASE approach, as delivered by Netskope, provides a number of design advantages over conventional enterprise networking:

Networking Benefits

- 1. Simplified operations:** Shifting the back-end networking to Netskope means that your organization only has to manage the tunnels to the NewEdge data center. This is further simplified by the Netskope Client, which can be used to negotiate the tunnel connections for users, as well as Borderless SD-WAN, which maintains and optimizes the connections for both users and offices automatically through policy.
- 2. High speeds:** By tapping into NewEdge, your organization has access to unrivaled cloud connectivity to a highly performant network, without the backhaul to the corporate data center for egress.
- 3. More cloud application coverage:** Most organizations maintain a small number of cloud interconnects. For example, an AWS shop might acquire Direct Connect to support the AWS development team, but rely on unpredictable open internet for other clouds or SaaS apps. Using Netskope, you can support a broader range of applications with speed and security.
- 4. More geographic coverage:** Internet service around the world remains highly fragmented, with unpredictable service quality and speeds from the vendors in various regions. Organizations can generally improve their overall global connectivity via a first hop through NewEdge.
- 5. Single-vendor consolidation:** Every region has different networking service providers, which creates a painful burden on the typical enterprise. Even within a region, service quality varies especially in diverse continents such as Europe and South America. By leveraging Netskope, organizations can ensure a high-quality user experience around the world without having to maintain multiple vendor relationships in each region.
- 6. Better security:** Instead of relying on a mixed set of network appliances and exceptions to policy, organizations can improve their security by using Netskope. Netskope SASE delivers protections that include threat protection, data protection, and zero trust network access as services. Shifting protections through Netskope helps the organization simplify and reduce the attack surface within their own networks, as illustrated in this guide.

INDEX

¹ "CIOs, CTOs and technology leaders: Latest findings from PwC's Pulse Survey," PwC, January 27, 2022.

² "Gartner, Hype Cycle™ for Cloud Security, 2021," by Tom Croll, Jay Heiser, July 27, 2021.

³ "Cloudy With a Chance of Malice," Netskope, February 23, 2021.

⁴ "Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2025," Statista, March 18, 2022.

⁵ "40% of organizations have suffered a cloud-based data breach," Security Magazine, October 29, 2021.

⁶ "A Stanford Economist Who Studies Remote Work Says Half of All Workers Will Make This Big Change in 2022," Inc., January 8, 2022.

