netskope

WHITE PAPER

Blueprint for Zero Trust in a SASE Architecture

Continuous Adaptive Trust—The Key to Adopting Zero Trust and SASE and How to Get There

For a concept that represents *nothing*, "zero trust" is absolutely *everywhere*. Lost in much of the noise, though, is a cogent description of the goals of a zero trust approach and the outcomes that approach is meant to achieve. Companies who've explored how to embark upon zero trust projects encounter daunting challenges: inconsistent definitions, incompatible and complex legacy IT, and indeterminate signs of completion. The shift to cloud, the transition to remote work, and digital transformation create an opportunity to frame the importance of zero trust principles unique to multi-cloud secure access service edge (SASE)-enabled architectures. At Netskope, we believe zero trust in a SASE or hybrid environment consists of the ability to establish continuous adaptive trust across users, devices, networks, applications, and data to increase confidence in policy enforcement everywhere.

The primary goal of a zero trust approach is to shift from *trust but verify* to *verify then trust*. Resources no longer place implicit trust (IP address, for example) in any entity that wants to connect. Through evaluation of several contextual elements—user identity, device identity and security posture, time of day, geolocation, business role, sensitivity level of the data, and more—the resource itself can determine *an appropriate level of confidence,* or trust, only for that specific interaction and only for that specific resource. As an example, the level of confidence of a user on a device with an agent capable of reporting rich telemetry reflecting the device's environment is likely to be greater than the confidence of a user whose device reports nothing other than a previously noted MAC address. Greater confidence confers higher trust.

But evaluating the confidence level at the start of any interaction is insufficient. During the interaction, context should be *continuously* evaluated. Alterations to the context can result in an *adaptation* (an increase or a decrease) in the level of trust, which is likely in turn to alter the type of access to the resource.

Binary decisions such as full access or no access lack the flexibility required by contemporary and emerging work styles—for example, a higher-risk SaaS application necessary for employee productivity but unmanageable with simple "all" or "nothing" access. Access should be context-aware, balancing trust against risk. In higher-risk situations, access

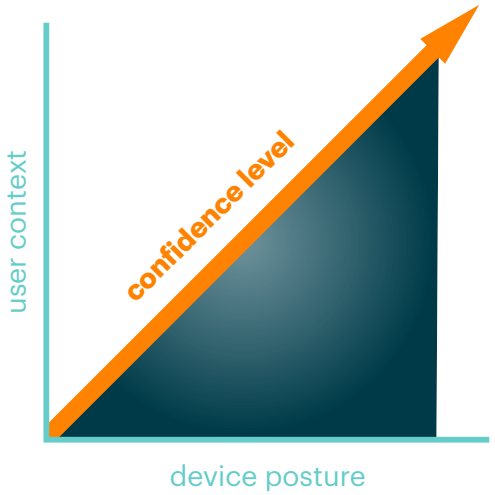

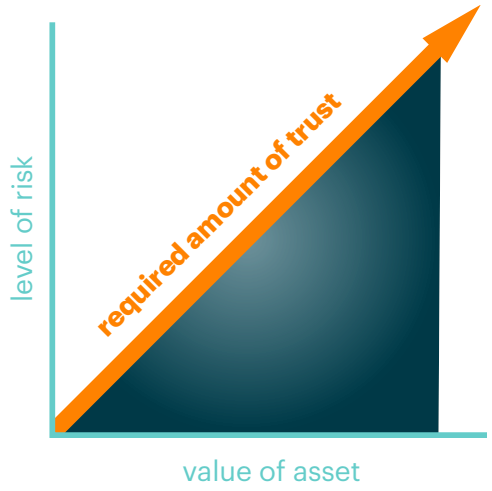

More importantly, though, a true zero trust approach reduces risk and increases business agility by eliminating implicit trust and by continuously assessing user and device confidence based on identity, adaptive access, and comprehensive analytics.

70%

of users continue to work remotely **through the first half of 2021.**

More than half of web gateway traffic is related to the cloud

APPS AND CLOUD SERVICES

90% of traffic is TLS encrypted*

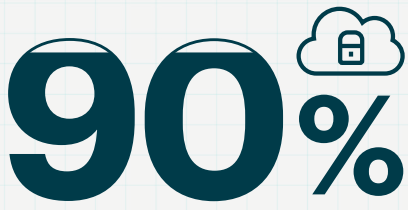

TLS ENCRYPTED

*Google HTTPS Transparency

is limited but not blocked completely, allowing some work to commence. In low-risk situations, access is expanded and certain obligations (such as MFA or a managed device) can be reduced. The continuous adaptive trust model increases the requirements for the level of confidence in parallel with the value of the asset being accessed. Based on signals such as application or activity risk, user risk, data sensitivity, device posture, user location, and other attributes, adaptive access provides the ability to make real-time decisions to permit, deny, restrict, redirect, and even coach the user. Adaptive access aligns policies with risk appetite, which may include revocation of access, if required at a point in time.

A secondary goal of a zero trust approach is to assume that the environment can be breached at any time, or even that it already has been, and design backward from there. Mostly a new mindset, such an assumption facilitates the deployment of patterns and practices that shrink attack surfaces, limit blast radiuses, constrain lateral movement, and respond to threats with greater speed and accuracy.

At its core, zero trust is more than just a restatement of two common but important security principles: least-privileged access and resource concealment. It's true that a zero trust approach results in approved users receiving only the necessary access to approved resources and that unapproved resources are invisible to and inaccessible by unapproved users. More importantly, though, a true zero trust approach reduces risk and increases business agility by eliminating implicit trust and by continuously assessing user and device confidence based on identity, adaptive access, and comprehensive analytics.

WHY ZERO TRUST MATTERS NOW

The traditional networking model originated in the era of workers in offices accessing corporate applications running from on-premises data centers. Companies originally built their network security in a manner similar to physical security, namely by presuming that bad guys were on the outside and only good guys were on the inside. With a decent level of protection against inbound threats, most companies found the newer concept of zero trust networking to be an object of curiosity, but nothing more. The perimeter was good enough for the time being; IT and security teams had other pressing concerns.

And then the world changed: Applications and users and data migrated outside the perimeter. Attracted to a wide array of options and their ease of procurement, companies began turning to Software-as-a-Service (SaaS) applications for an increasing number of business processes. Indeed, the average enterprise with 500-2,000 users now subscribes to 805 *distinct SaaS applications!*

Source: Netskope Cloud and Threat Report - July 2021

Of course, some business processes still required custom applications. Companies discovered that Infrastructure-as-a-Service (IaaS) and Platform-as-a-Service (PaaS) public clouds offered greater agility than traditional data centers; these became the default target for most new applications. Indeed, it's now the case that most companies embrace a multi-cloud strategy, choosing from multiple hyperscale providers based on application requirements and developer proficiency. (Lift-and-shift, the process of migrating existing applications to the public cloud, has been less successful, not least because the cloud amplifies the problems of bad architectures and bad security decisions.)

2020 demonstrated that companies can adapt to massive change when much of the world's workforce suddenly transitioned to remote work. The change was initially disruptive, and not all industries can or should sustain a mostly remote work style. But for many companies, remote work is already the new normal. Employees and their business partners require access to applications wherever they reside (on premises, SaaS, public cloud) from whatever kind of device (managed or unmanaged).

Yesterday Versus Today

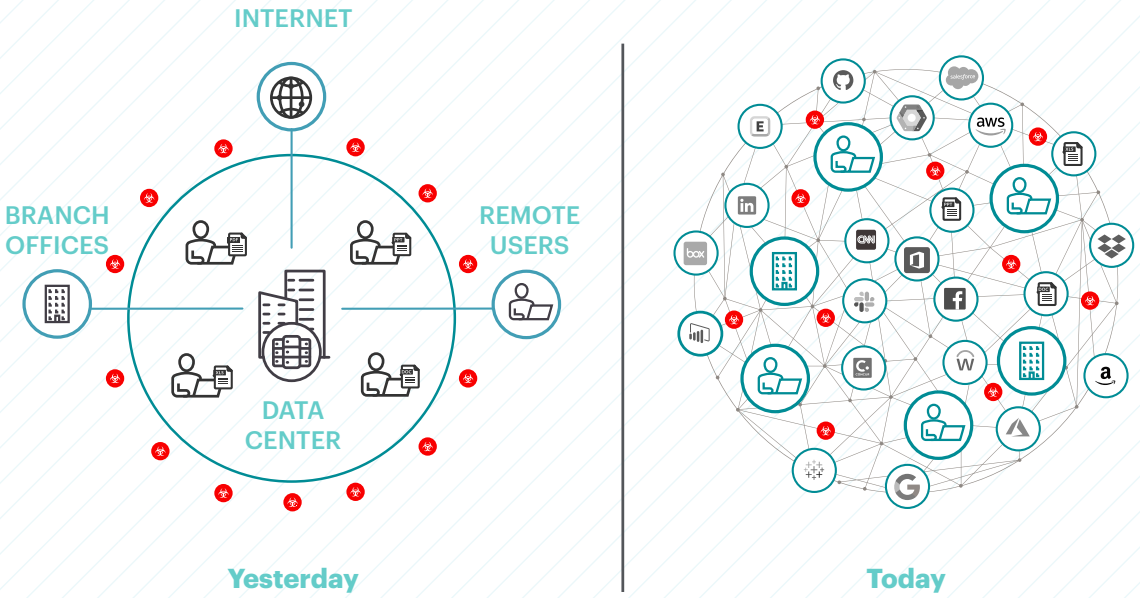

The zero trust model is ideally suited to accommodate such a requirement. Broadly, two categories of products implementing this model have emerged: zero trust network access (ZTNA) and identity-based segmentation (microsegmentation).

Products in the ZTNA market apply zero trust principles to the process of making applications available to users. A zero trust–based approach to private applications, whether on premises or in the public cloud, and to SaaS applications can be granted to specific users regardless of what network they might connect to. The degree of access is adapted to the context surrounding the interaction, which in turn provides a level of confidence commensurate with the validity of the session. For example, users on unmanaged devices can be granted full access to managed applications processing public data, read-only access to managed applications processing sensitive data, and no access to applications processing confidential data. Furthermore, analytics that combine historical user behavior with an understanding of specific application functionality can detect and mitigate threats before they wreak havoc.

Identity-based segmentation (as defined by Gartner) is a new way of looking at microsegmentation. Microsegmentation is a form of isolation enforced by a static set of rules that define which resources (e.g., a device, workload, or container) can communicate with each other. Identity-based segmentation expands on this concept, adding dynamic rules that assess identity and other attributes as part of determining whether to permit access. Resource identities are portable and independent of the underlying network. Segmentation based on identity plus additional context measured in real time enables greater flexibility, agility, and extensibility of segmentation control and policies.

Of course, it must be acknowledged that zero trust models necessarily add a degree of management overhead. Owners of resources must assume the responsibility of carefully assessing and continuously adjusting the lists of allowed users for their resources—a process known as entitlement management. Management of entitlements can be a mostly manual process, but new emerging technologies seek to provide a welcome and beneficial automation that can reduce the errors normally accompanying any manual process. Furthermore, a zero trust approach requires managing not only entitlements but also defining the attributes and contextual elements that together determine the level of confidence required for interacting with resources. In fact, in environments where traditional access and authorization models can't be well-defined, the addition of context enables a more accurate assessment of trust.

INTRODUCING THE SASE ARCHITECTURE

Security's origins began in the network. Over time, numerous tools jockeyed for position in data centers and competed for attention from administrators. They were generally effective when applications and data remained on premises and users worked from traditional offices. Some of the tools offered mechanisms to communicate amongst themselves. But as users moved out of offices and data and applications moved into the cloud, legacy tools became blind. They all operated under one overarching assumption: Applications, data, and users are static. Because this assumption is no longer valid, those tools have lost much of their suitability. They don't work with each other, they don't scale well, they lack unified administration, and—crucially—they can't perform their functions when your data is stored and processed on someone else's infrastructure.

SASE, an architecture defined by Gartner in 2019, promises to overcome the limitations of too many tools and too many consoles. SASE combines common network functions (e.g., SD-WAN, WAN optimization, QoS, routing, CDN, others) with common security functions (e.g., SWG, CASB, ZTNA, VPN, FWaaS, RBI, others) into a consolidated architecture with unified administration. Policies apply access control to any application or service and monitor and control the movement of sensitive information from and to all users and resources. SASE delivers network and security functions from the cloud, ensuring a consistent user experience wherever users and applications reside.

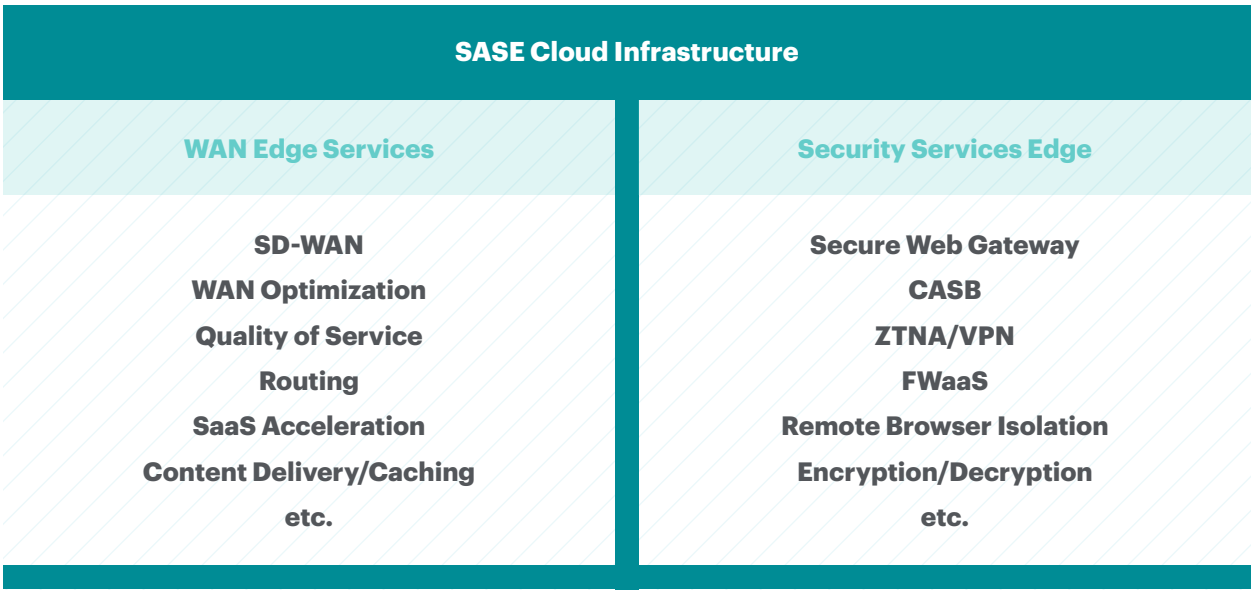

Source: Gartner, "2021 Strategic Roadmap for SASE Convergence," G00741491, https://www.gartner.com/document/3999828

Good SASE architectures implement zero trust principles. SASE consolidates all the ways users access resources, while remaining neutral to how confidence is assessed and access is granted. Zero trust principles insist that access is granted and confidence is monitored based on sets of conditions, while remaining neutral to any given technical architecture. When combined, SASE and zero trust represent a fundamental change to the ways companies protect their digital assets. In fact, SASE can be a good basis from which to develop an effective zero trust program that encompasses fully hybrid environments in which users, applications, and data can be anywhere.

At a high level, and depending on vendor capabilities, companies who adopt a SASE architecture with zero trust principles can expect to:

- Gain user risk and application risk insights to determine a level of confidence in the access granted under varying conditions and adapt access based on the confidence factor.
- Extend zero trust principles beyond only private applications to web and SaaS applications based on risk insights with adaptive policies and postures.
- Apply risk insights within applications to control access to specific activities (for example, a low-trust scenario that permits viewing and commenting but prohibits sharing and deleting).
- Activate additional security services such as remote browser isolation and advanced data loss prevention based on assessed level of risk or trust.
- Continuously monitor for changes in context that require reassessment of trust (for example, reauthentication, step-up authentication, alteration of permissions, or increased/decreased access).
- Reduce overall attack surface area by eliminating the exposure of protocols and services to the public internet.

OUTCOMES FOR THE BUSINESS AND THE USERS

Modern digital business will not wait for permission. At the same time, modern digital business increasingly relies on applications and data delivered over the internet (which, in fact, wasn't designed with security in mind). Companies can achieve their business goals without making security tradeoffs by building continuous adaptive trust principles into their security and risk programs and into their digital transformation plans from the outset.

Agility in the business

Agility in business requires elasticity in infrastructure and services both in scale of volume and location and in breadth of new services and applications. Outcomes that can and should be achieved based on zero trust principles in a SASE and hybrid architecture include:

- Harmonized user experience—access is the same regardless of where a user is; access is predictable regardless of the user's device type.
- Location independence—application access is decoupled from the underlying network design; applications can be moved (say from on premises to the public cloud) without forcing users to change their habits.
- More opportunities to provide some degree of access—to reorient the majority of security decisions away from "no" toward "yes, with conditions."
- Increased collaboration with suppliers and partners without provisioning local user accounts for them and without placing demands on their computing environment.
- Proactive security operations built to support the growth of applications and to eliminate the fire drills of hunting down and retroactively securing access to applications after they're already deployed.

Reduction of risk

Cyber risk is a priority for most boards of directors. Every company must determine its own risk appetite and tolerance; these tend to be roughly the same for most companies in any given industry. Managing risk is complicated by dependency on increasingly lengthy and opaque supply chains, the proliferation of cloud services and applications, and an ambiguous regulatory environment. Risk reduction in this new environment will involve a zero trust approach in which:

- Resources are moved from public to private access and thus shielded from the internet, invisible to those who lack strong access credentials or the ability to demonstrate confidence.
- Inappropriate access is constrained, reducing the blast radius of compromised accounts.
- Visibility into sensitive data types, locations, and movements is improved and constant.
- Analytics offers the whole picture of acceptable policy and behavior, thereby more rapidly surfacing risks and threats (both anomalous and malicious activity) for quick containment and neutralization.
- Overall security posture is improved, making companies less attractive to attackers.

Streamlined product deployment and maintenance process

Business agility and risk reduction require the right architecture, including:

- A cloud-native security service that scales as the business requires, eliminating the deployment complexities and capacity constraints commonly associated with hardware security appliances.
- A single cloud platform, single-pass inspection, and policy enforcement point, driven by a single console and policy engine that is applied to ensure a consistent security policy across all channels.
- A single vendor relationship that eliminates delays commonly accompanying troubleshooting and repairing products with untested or unknown interoperability characteristics.

FIVE PHASES FOR CONTINUOUS ADAPTIVE TRUST WITH NETSKOPE

Netskope helps companies achieve their continuous adaptive trust goals. The journey might not be exactly the same for every company. In fact, your industry and your degree of cloud adoption and your existing legacy systems will likely influence not only where you begin your journey but also the amount of effort necessary.

A few simple steps can help every company prepare. Start with these tasks:

- Devise a set of user personas that reflect typical business roles. List common access requirements for each persona. Start from scratch—don't rely on existing role-based access definitions.
- Inventory all private applications and SaaS applications. Create a map illustrating how applications and their components interact with each other and with resources external to the company.
- Identify all data assets: their locations, their sensitivity levels, their business functions, and their intended lifetimes.

Every successful continuous adaptive trust project demands a robust identity and access management program. Without an accurate and reliable system of record for identity, a high level of trust can't be determined or inferred. Every entity—whether human, device, or object—must present an identity that every other entity can validate. Fortunately, most companies already have some kind of identity management system in place. Crucially for zero trust, the system must be compatible with common standards like SAML, OIDC, and OAuth because most organizations expect this. Identity federation establishes a degree of trust between disparate identity realms, allowing users in one realm to access authorized resources in other realms without possessing separate identities.

With this information in hand, you can begin planning and executing the five phases on the following page.

PHASE 0 — Before You Begin

- Devise personas and assign access requirements
- Inventory all SaaS and private applications
- Inventory all data assets

PHASE 1 — Zero Trust Access

- Define sources of truth for identity
- Construct a mapping of users to applications
- Remove stale or unused entitlements
- Steer all access through a policy enforcement point to remove risks of direct access

Establish this zero trust access baseline

PHASE 2 — Adaptive Access

- Add context to access policies
- Allow context to signal whether to require strong authentication
- Incorporate context into coaching
- Continuously adjust access policies

Enrich the trust baseline for authorization within app activities

PHASE 3 — On-demand Isolation

- Insert remote browser isolation for access to risky and low-reputation web sites
- Monitor for command-and-control attempts and other anomalous behavior

Apply explicit trust controls to destinations

PHASE 4 — Continuous Data Protection

- Differentiate access by managed and unmanaged devices
- Adapt access policies based on context
- Ensure correct configuration of cloud resources
- Identify sensitive information and control its spread
- Continuously assess sharing permissions and app-to-app integrations

Continuously investigate and remove excess trust. Adopt and enforce a least-privilege model everywhere

PHASE 5 — Inform & Refine With Real-time Analytics

- Maintain visibility over applications and risk
- Inspect cloud and web activity for best assessments of and adjustments to policies
- Create visualizations tailored to various stakeholders
- Implement closed-loop policy refinement

Strengthen security and trust posture with closed loop refinement of policy

PHASE 1

Zero trust access: don't allow anonymous access to anything

Begin building the context from the point of first access. Start with shoring up identity and access management (including roles and role membership), private application discovery, and a list of approved SaaS applications and website categories. Reduce the opportunities for lateral movement and conceal applications from being fingerprinted, port scanned, or probed for vulnerabilities. Require single sign-on (SSO) with multi-factor authentication (MFA).

The Netskope Security Cloud enables federated SSO with adaptive access controls to websites, private applications, and thousands of SaaS applications, including shadow IT adopted by various business units. Netskope Private Access delivers full ZTNA capabilities from private application discovery to application level access for managed users or third-party contractors or M&A use cases. The Netskope Next Generation Secure Web Gateway (NG SWG) as an access control point to web and SaaS can thwart phishing and other attacks against credentials for SaaS, webmail, and cloud storage, which represent the top threats most companies face today.

Specific tasks for this phase:

- Define the source of truth for identity and what other identity sources they might federate with.
- Establish when strong authentication is required.
- Construct and maintain a database that maps users (employees and third parties) to applications. This necessitates regular conversations with business units. Nominate a specific set of individuals to hold this responsibility.
- Rationalize application access by removing stale entitlements (of employees and third parties) that are no longer required because of role changes, departures, contract terminations, etc.
- Remove direct connectivity by steering all access through a policy enforcement point for every private or internal application (ZTNA) and for access to SaaS and the web (CASB, SWG).
- Maintain real-time application and user dashboard and visuals. Control which users should have access to which apps and services.

PHASE 2

Adaptive access: maintain the explicit trust model

Add more context so that access control becomes adaptive to maintain an explicit trust model. Evaluate signals from application risk assessments, user risk assessments, endpoint posture checks, and locations of users, applications, and data. Implement adaptive policies that invoke step-up authentication, raise an alert for the user to indicate whether to proceed with or cancel an application activity, provide a business justification to continue, or real-time coach toward approved applications.

The Netskope Security Cloud provides deep context across inspected web and cloud traffic and delivers real-time application risk scores and insights, user-risk scoring, and UEBA insights over time based on user behavior. These scores activate real-time adaptive policy actions such as step-up authentication, process termination, and more, based on the desired policies for users, data, or applications.

Specific tasks for this phase:

- Determine how to identify whether a device is managed.
- Add context to access policies (block, read-only, or allow specific activities depending on various conditions).
- Increase use of strong authentication when environmental risk is high (e.g., for all remote access to private apps and SaaS to delete content) and decrease its use when risk is low (e.g., managed devices accessing local applications for read-only).
- Evaluate user risk; coach classes of users toward specific application categories.
- Continuously adjust policies to reflect changing business requirements as applications evolve, as new ones arrive, and as old ones are decommissioned.
- Establish a trust baseline for authorization within app activities.

PHASE 3

On-demand isolation: contain the blast radius

Decouple browsers from applications and place controls on outbound traffic. Deploy remote browser isolation (RBI) to minimize browser functionality to that of a display device. Enable an on-device firewall to constrain the destinations to which the device can establish connections.

The Netskope RBI is native, single-pass, and integrated with existing data protection and threat neutralization capabilities. Plus, it offers a degree of performance not available from RBI products service-chained to traditional proxies or network gateways. In keeping with the theme of removing implicit trust, direct access to risky web resources should be minimized, especially as users simultaneously interact with managed applications.

On-demand isolation—that is, isolation that automatically inserts itself during conditions of high risk—constrains the blast radius of compromised users and of dangerous or risky websites. Furthermore, the Netskope Cloud Firewall is a highly effective mechanism that thwarts attempts to communicate with attacker-owned command-and-control nodes.

Specific tasks for this phase:

- Automatically insert remote browser isolation into the path for access to risky websites with poor reputations.
- Configure remote browser isolation for access to private applications from unmanaged devices.
- Evaluate remote browser isolation as an alternative to CASB reverse proxy for SaaS applications that behave incorrectly when URLs are rewritten.
- Monitor real-time threat and user dashboards for command-and-control attempts and anomaly detection.

PHASE 4

Continuous data protection: apply single-pass data policies

Gain visibility into where sensitive data is stored and where it spreads. Monitor and control movement of sensitive information through approved and unapproved SaaS applications, websites, public cloud storage, custom applications in public clouds, and outbound email. Apply single-pass data protection policies across inspected web, email, and cloud traffic for data in motion and via API for data at rest.

The Netskope Security Cloud provides control over unintentional and unapproved data movement plus full analytics of all inspected data activity. Legacy defenses are unable to decode cloud traffic and thus are unable to apply the necessary data protection, especially for personal instances of approved and unapproved applications. The Netskope Cloud Firewall-as-a-Service (FWaaS) provides outbound network access control of all ports and protocols for remote users and branch offices, mitigating command-and-control attacks and data exfiltration.

Specific tasks for this phase:

- Define overall differentiation for data access from managed and unmanaged devices.
- Add adaptive policy details to access content based on context (e.g., public content full access on unmanaged devices, sensitive content read-only on unmanaged devices, confidential content blocked on unmanaged devices).
- Invoke cloud security posture management (CSPM) to continuously assess public cloud service configurations to protect data and meet compliance regulations.
- Assess the use of inline DLP rules and policies for web, managed SaaS, shadow IT, public cloud services, and email to protect data and meet compliance regulations.
- Define data-at-rest DLP rules and policies for managed SaaS and IaaS environments, especially file-sharing permissions for cloud storage objects and application-to-application integrations enabling data sharing and movement.
- Fine tune policies by adding user or group attributes when applicable.
- Continuously investigate and remove excess trust. Adopt and enforce a least-privilege model everywhere.

PHASE 5

Inform and refine with real-time analytics and visualization: maintain a strong security and trust posture

Enrich and refine policies in real time. Assess the suitability of existing policy effectiveness based on user trends, access anomalies, alterations to applications, and changes in the sensitivity level of data. Adjust policies accordingly to tolerate exposure to risk without drifting out of bounds.

Netskope real-time analytics and visualization provide insights into operations across the Netskope Security Cloud, identifying who should (and who shouldn't) be working with the company's data in a given application set. These capabilities work in conjunction with security operations, network operations, and threat hunting teams, and explain the results of the security program to executive staff and application stakeholders. Insights derived from analytics help organizations plan and anticipate where their security program is going and what it must do next.

Specific tasks for this phase:

- Maintain visibility into the applications and services the company uses, and the associated levels of risk.
- Find out who should and shouldn't have access to private resources, and dial back access to minimal levels to reduce exposure.
- Gain greater visibility and establish a deep understanding of cloud and web activity for ongoing adjustments and monitoring of data and threat policies.
- Identify key stakeholders for your security and risk management program (CISO/CIO, legal, CFO, SecOps, etc.) and apply visualizations to the data that they can understand.
- Create the dashboards to get visibility into these components: site, application, instance, user, activity, file, source/destination, and more.
- Ensure the ability to import and export dashboards to collaborate with other security teams.
- Strengthen security and trust posture with closed-loop refinement of policy.

Digital transformation has been accelerated by the pandemic events of 2020, causing businesses to evaluate how to re-architect their network and security infrastructure and programs to meet current needs. It's clear a new approach is required to enable a fast, easy user experience with simple, effective risk management controls across multi-cloud and hybrid architectures.

Netskope safely and quickly connects users directly to the internet, any application, and their infrastructure from any device, on or off the network. With CASB, SWG, and ZTNA built natively in a single platform, the Netskope Security Cloud provides the most granular context, via patented technology, to enable conditional access and user awareness while enforcing zero trust principles across data protection and threat prevention everywhere. Netskope's global security private cloud provides full compute capabilities at the edge, to deliver strong security, high performance, and reliable global networking.