netskope

White Paper

Zero-Trust Security Model Applied to Netskope Intelligent SSE

EXECUTIVE SUMMARY

The current situation with legacy on-premises security defenses trying to support a hybrid work environment and zero-trust principles is challenging for companies. Complications can include poor user experience, complexity of disjointed solutions, high cost of operations, and increased security risks with potential data exposure. Simple “allow and deny” controls lack an understanding of transactional risk to adapt policy controls and don’t provide real-time coaching to users. The implication of continuing down this path is poor use of resources, limited business initiatives, frustrated users, and an increased potential for data exposure and regulatory fines. The binary option of blocking new applications and cloud services impedes digital transformation and innovation, and an open allow policy lacks data protection for transactional risk and real-time coaching to users on risky activities.

This white paper takes the position of consolidating security defenses into a security service edge (SSE) cloud platform to reduce security risks and protect data, support zero-trust principles, improve user experiences, and to enable business agility. This position supports a wider range of use cases with adaptive access policy controls related to transactional risk while providing least-privilege access and protecting data. As a result, I&O leaders should apply zero-trust principles to SSE for an understanding of how together they open new use cases often required for hybrid work environments and adopting SaaS. The benefits include reducing security risks and protecting data, reducing cost of operations, freeing up full-time employees for new projects, and increasing business agility.

WHAT IS DRIVING THE CHANGE TO ZERO-TRUST PRINCIPLES?

Perimeter security is no longer effective and legacy security approaches lack empathy for the user experience and disregard digital productivity and its impact on business efficiency. Disjointed networking and security teams using an onsite hardware-centric approach are struggling to meet new requirements for cloud transformation and to support hybrid working. As cloud-based applications and data increasingly sit outside the data center castle, and as the pandemic has pushed employees outside the LAN moat and its controls, commonly siloed networking and security teams must cooperate to realize the benefits of zero-trust principles.

Furthermore, these principles have evolved. Network segmentation using hardware-centric solutions was once the focus; however, migration to the cloud, hybrid working, and SaaS applications impact its ability to scale, making this legacy approach cost prohibitive and impractical. With the dissolving network perimeter, data being hosted across multiple cloud and on-premises locations, and employees preferring the flexibility of hybrid work, there is an urgent need to shift toward cloud-centric security solutions that can employ a zero-trust approach to follow and secure the data wherever it goes, while seamlessly scaling to support digital transformation initiatives. Today, zero-trust principles offer just the right access to just the right resources by just the right people at just the right times for just the right reasons — as defined by context in risk-based policies.

WHAT IS THE ZERO-TRUST MODEL?

Trust in digital systems is an old concept that far predates Forrester Research's popularization of doing away with it. Forrester's notion of removing all trust contains three principles. First, all entities are untrusted by default and risk should be assessed for every session. Known devices and known networks no longer implicitly trust each other, and users no longer retain unconditional access. Second, least-privilege access provides the bare minimum for each transaction, allowing access to only approved resources and concealing the rest. Here we lean into context and content being the new perimeter to determine what defines risk-appropriate access for the current session. Third, comprehensive security monitoring closely observes user and asset activities, behaviors, and trends to find unknowns and refine least-privilege access. Rich context for users, apps, risks, and data beyond what legacy onsite networking security solution logs have provided is required to support zero-trust principles.

So, the zero-trust security model is not something you buy, or a buzzword misaligned in marketing content, nor is it just networking segmentation, identity services, or a security awareness training program. It's an essential business strategy aligned with a zero-trust security model.

Forrester further defines seven operational domains — five for security controls and two for cross-domain interactions:

- **Data & Assets** — to secure, manage, categorize, and classify data, plus encrypt data at rest and in motion
- **Users & Identities** — enforcing user access and securing users on the internet for their identities
- **Devices & IoT** — the ability to isolate, secure, control, and remove any networked device
- **Applications & Workloads** — front-end and back-end systems, SaaS, and IaaS requiring zero-trust principles including least-privilege access and continuous monitoring
- **Networks** — the ability to segment, isolate, and control networks
- **Visibility & Analytics** — to understand trends, behaviors, anomalies, and find unknowns through rich context
- **Automation & Orchestration** — to integrate security solutions, reduce response times, and increase analyst capacity

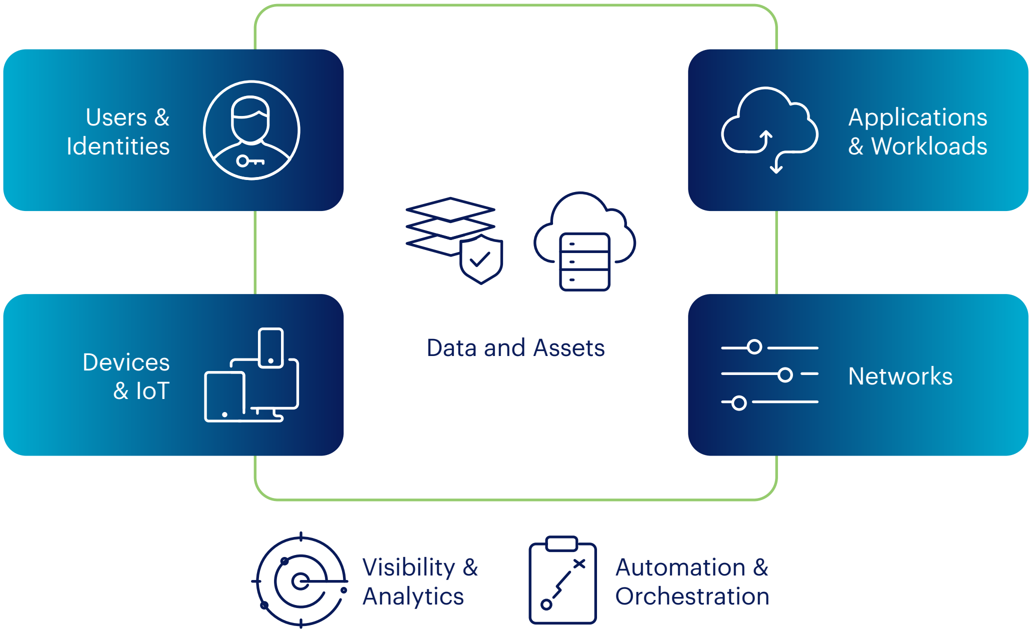

Figure 1: Zero-trust security model with five core components and two process and technology competencies.

The zero-trust security model provides the opportunity to review your processes and technologies to introduce zero-trust principles, assess the current expense-in-depth security model, and reduce vendor sprawl to consolidate, reduce complexity, and lower total cost of operations. To accelerate your journey, plan to work across your organization and assess your current technologies for zero-trust principles, plus align with your business initiatives including hybrid work, cloud transformation, innovation projects, and compliance regulations.

WHAT LINKS EVERYTHING TOGETHER FOR ZERO TRUST?

At a surface level reviewing zero-trust reports and diagrams, you will see data often at the center, and with our legacy perimeter security mindsets we mentally put a moat and castle around it. However, after certification from the Forrester Adopting Zero Trust certification course, you learn why data is at the center, because it flows through, resides, or touches all the other components in the security model. Data flows like a river where a perimeter security model — assuming data resides only with managed devices, networks, applications, data centers, and users — can no longer protect it. In some cases, data is the only element you directly own and control as business partners use bring your own devices (BYOD) to access third-party SaaS applications from networks beyond your control. Now you are probably thinking here comes the data loss prevention (DLP) marketing pitch; however, zero-trust principles provide data protection before DLP is invoked using least-privilege access and context to assess risk for each transaction. This makes the use of DLP more focused, efficient, and effective when applied with zero-trust principles.

WHAT ARE THE DATA SECURITY INTERSECTION POINTS FOR THE ZERO-TRUST COMPONENTS?

There are four intersection points in the zero-trust model for the data component where security teams need to understand how their organization collects, transmits, stores, and shares data. Data can be transitory and dynamic through its life cycle where understanding what constitutes data, where it resides, along with its value and classification are important. Data intersects with users and identities, devices, applications and workloads, and networks, as shown in Figure 1 above. Just as important is how data is accessed, utilized, moved, and shared using visibility and analytics to find the patterns, behaviors, and anomalies, plus the unknown risks.

User and identities require authentication and authorization to the other zero-trust security model components for least-privilege access. What applications can be accessed, across what networks, for what devices, and most importantly for what data and entitlements are key use cases to understand. Identity access management (IAM) with identity services for single sign-on (SSO) and multi-factor authentication (MFA) are critical controls to mature for this intersection. While passwords may not be entirely retired, zero trust seeks to reduce their use when possible; expect to see security vendors hasten their demise. More advanced use cases may use privilege access management (PAM) for administrators, role-based access controls for separation of duties, and step-up authentications in real time based on the risk of a transaction.

Users consist of employees, contractors, business partners, customers, and the public. The zero-trust model requires controls between these users and assets with no implicit trust and least-privilege access. The ability to distinguish normal from abnormal user behavior is becoming increasingly important to detect insider threats, malicious external actors, abuses of access and entitlements and the often-associated data theft and exfiltration. User and entity behavior analytics (UEBA) with an understanding of rich context of users, devices, apps, instances, activity, network location, and data sensitivity is required. For example, it is common to see an increase in data exfiltration from company instances to personal instances of applications within the last 30 days of employment before users depart. For applications that do not provide instance awareness for users, the ability to map company and personal logins with identity mapping extends this use case.

Beyond UEBA machine learning and detectors is the ability to provide a user risk score, or confidence index for their activities, events, alerts, and normal versus abnormal behaviors. User risk scores can then be used to identify potential insider threats, risky users, and be applied in real time for adaptive access control policies. Consider account compromise for a user with multiple failed logins and proximity alerts, plus higher than normal data activity. Policy controls should be able to automatically assign the user to a high-risk group, restrict access, and share the user risk score with other security solutions. Here, adaptive access policy controls during business transactions and workflow automation with integration to IT service management and collaboration solutions are critical to protect your company, assets, brand, and reputation.

Devices are both managed and unmanaged, and beyond laptops, tablets, smartphones, and watches also include the internet of things (IoT) such as smart cameras, smart lighting, printers, sensors and monitors, and networking devices as examples. A work-from-anywhere environment also includes BYOD and virtual desktop infrastructure (VDI), all realities of less control for endpoints. For managed devices, you likely already use an endpoint detection and response (EDR) solution and invoke disk encryption to secure data at rest, plus run regular data backups. EDR paired with an extended detection and response (XDR) platform or a managed detection and response (MDR) service is also popular to prevent, detect, and respond.

Device health checks, device risk scores, or device confidence ratings should also be assessed before access is granted to networks, applications, and data for transactions within a zero-trust security model. Never assume a default level of trust for devices, and this perspective also makes securing IoT easier. Inventory and characterize existing and new devices and IoT on your network with a device security posture and risk rating when possible. And as we noted for user risk scores, the same holds true for device risk using it within real-time adaptive access control policies and in analytics to find the unknown risks. From a network perspective, segment and isolate users and devices away from the rest of the network to minimize the impact of a cybersecurity incident. For the device component of the zero-trust security model, you need the ability to isolate, secure, control, and remove any connected device.

Applications and workloads — The adoption of SaaS applications and the utilization of IaaS, along with mobility and remote working, have had the most impact to retire perimeter security models and passwords. The average mid-sized company utilizes over 800 applications while a larger enterprise can use 2,400 or more, based on Netskope research. Less than 3% of these applications are managed; the rest are freely adopted by business units and users as they march toward digital transformation, also known as Shadow IT or more appropriately Business IT. Expect SaaS adoption to grow given the forecast to increase nearly 20 percent year over year for the next five years. The leading candidates are office productivity suites, customer relationship management (CRM), and cloud storage, plus the strong interest in generative AI with apps like ChatGPT or CoPilot. A key point to remember is the personal side of SaaS and mobility use that blends into a work/life balance, where data can easily flow between a company instance of cloud storage and a personal instance for the same application.

As data increasingly moves to cloud-based SaaS and IaaS within this transformation, it requires inspection inline and at rest with zero-trust principles. Like users and devices, applications have risk ratings for use in adaptive access policy controls, guiding users to use safer application alternatives, and as part of a third-party risk assessment. Allow and deny controls with firewalls and web gateways are challenged with the quickly changing environment for applications. To deny access to SaaS applications where more than 97% are unmanaged, you deny the digital transformation journey for your business success. If you allow without the ability to inspect inline application activity, instance, risks, and data sensitivity, you leave the door open for insiders, data exfiltration, and theft.

Just as firewalls inspect packets across ports and protocols, and web gateways inspect HTTP/S web sessions, SSE solutions with cloud access security broker (CASB) capabilities decode applications inline to understand the context and content of a transaction for an adaptive access policy control decision based upon zero-trust principles. This same decoded rich context viewed through analytics also uncovers unknown risks to further refine least-privilege access to applications and data in a closed loop. Context and content is the new perimeter in a hybrid working environment with a larger breadth of users and devices having direct access over the internet to applications from any location or network.

For in-house built applications, or DevOps, security teams need to become part of the pipeline from start to finish with gates and security checks to create an integrated DevSecOps team. Traditional security and development teams have been at odds; however, the zero-trust security model moves security into a business enabler and initiative alongside development in a partnership for success. Security champions within development can see the advantages of the integration to reduce software vulnerabilities, software supply chain issues, and web application weaknesses. Like network segmentation, workloads can benefit from microservices as they invoke API services and from microsegmentation versus one large application.

The intersection of data with SaaS applications, IaaS, and in-house built applications is perhaps one of the more challenging transformations for a zero-trust security model; it requires a mindset shift. Where is my data flowing, within what apps and risk profiles, what users and devices and their risk profiles, and on what networks? Can my analytics and reporting show unknown and unapproved data movement, and for the same application for company and personal instances, and across hundreds of applications? When an employee departs, can my security analysts assess the last few months of data movement and application use? As SaaS applications regularly update, are new data paths and transactions detected, assessed, and secured? These are modern-world use cases for applications, and data security legacy perimeter defenses were not designed to address with appliances or cloud hosting.

Networks leverage well-known security controls in the zero-trust security model, including next-generation firewalls (NGFW), secure email gateways (SEG), secure web gateways (SWG), cloud access security brokers (CASB), zero-trust network access (ZTNA), with the last three integrated into security service edge (SSE) solutions. Data loss prevention (DLP) across these solutions and on endpoints is at the intersection of data with networks, plus the inline policy controls of CASB solutions for apps, activity, instance, risks, and data sensitivity handling most of the burden. Note managed applications can also benefit from CASB API inspection of data at rest and SaaS security posture management (SSPM).

These network defenses also provide advanced threat protection and likely include remote browser isolation (RBI) to protect users from risky websites, uncategorized sites, newly registered domains, newly observed domains, and parked domains. The addition of egress firewalls to SSE solutions expands protection and application controls for non-web traffic, plus provides domain name system (DNS) security and intrusion prevention system (IPS) defenses.

One of the more interesting shifts in network security is the migration from virtual private networks (VPNs) to ZTNA. From a zero-trust security model perspective, ZTNA is inside-out in design, and unlike a VPN service that is publicly exposed, ZTNA has very specific access to an application or resource, unlike a VPN that may provide excess lateral movement. Remote access compromise is one of the leading entry points for ransomware, advanced threats, and espionage. Essentially ransomware is monetizing weak remote access and showing the need for zero-trust network access and zero-trust principles.

Technically, one of the goals of ZTNA is to remove the need for servers and endpoints to listen for random traffic. A VPN solution marketed as ZTNA still requires IP address pools, assigning IP addresses to endpoints as the solution is connecting devices to networks. While ZTNA connects users to applications, there is no assigning of IP addresses or for endpoints and servers to listen and respond to ICMP requests, as this creates a security risk where ICMP can be abused. Using ZTNA to replace VPNs to secure RDP, SSH, and other remote access methods is increasingly popular with zero-trust initiatives. The new combination of ZTNA with a software-defined WAN client also enables voice and IT service access applications to endpoints that standalone ZTNA does not address. Like passwords, VPNs are dying in this transformation.

The zero-trust security model also includes the concept of network segmentation; while it was the core focus early on, it did not scale, and was cost prohibitive and impractical. Microsegmentation takes the concept even further into fine-grained controls of applications, user access, and data repositories. These projects can be challenging and should start small in non-critical environments. No network should be flat with open lateral movement for an attacker; you will have to balance out how much segmentation fits into your security budget as you assess the full zero-trust security model, its components, and your use cases.

HOW DOES ZERO TRUST RELATE TO BUSINESS INITIATIVES?

Today, even the most basic business processes are rarely within the four walls of an organization. The zero-trust security model is better adapted to digital transformation than a perimeter-centric security approach by focusing on data, applications, and users. Companies that have adopted zero-trust principles often exceed compliance regulations and find it easier to support new business and operational models. Security becomes a business initiative, not a set of technologies or an expense. Zero-trust principles start with business initiative strategies, objectives, and tactics to enable agility, build trust and confidence, and differentiate solutions and services from competitors. Today more than one-third of companies commercialize and share their data as a business initiative where it needs protection from advanced threats and data breaches.

WHERE DO YOU START FIRST AND HOW LONG IS THE JOURNEY?

People and processes come first rather than technology next to a great user experience where security is transparent with little to no friction. Start by mapping out business use cases and processes before assessing technologies. Leverage zero-trust principles to increase freedom and reduce frustration, building trust as a competitive advantage for your organization. Many companies start by securing remote workers for cloud and web access, plus to private applications and resources. The legacy approach of placing hardware devices in employees' homes is expensive and lacks scale, just as backhauling employee traffic to on-premises infrastructure creates bottlenecks for traffic and often results in a poor user experience.

The easier solution to secure remote workers is a software client on employee devices connecting to cloud-edge security services leveraging cloud performance and scale as required. The integration of cloud access security broker (CASB), security web gateway (SWG), and zero-trust network access (ZTNA) into a security service edge (SSE) cloud security platform with global on-ramps makes it possible. While ZTNA has zero trust in its name, do not confuse it with the larger scope of the zero-trust security model and zero-trust principles; it was a naming choice by an analyst firm years ago to represent a specific technology.

Another popular starting point is prioritizing and securing business application traffic. SaaS-based applications like Microsoft 365 and Salesforce need direct-to-internet connections for remote workers and branch offices. SSE inspection of web, SaaS, and IaaS user traffic for any user, device, or location follows the zero-trust security model to retain visibility and control across the entire digital business ecosystem. Choosing to bypass traffic inspection of office productivity suites at the risk of data theft, exfiltration, or advanced threat entry for user experience is a tradeoff you no longer need to accommodate.

WHO IS INVOLVED FOR A TRANSFORMATION TO ZERO TRUST?

Traditionally, security teams are often introverted working in a silo and holding a reputation for saying “no” when included in larger business initiatives. This changes with the zero-trust security model where security is extroverted working with networking, infrastructure, and business units to enable new initiatives with a “yes” mindset aligning zero trust with business goals. Security needs to accelerate revenue growth and not just avoid costs, earn and build customer trust, protect intellectual property and competitive advantages, plus protect the company brand. So, zero trust is much larger than just ZTNA for secure remote access; it is a company culture about security enabling the business and should be looked at with a business value assessment perspective for top use cases. Effective security also embraces — even improves — application and network performance: This can be a common goal that catalyzes and unites SecOps, I&O, and NetOps teams.

HOW DOES NETSKOPE INTELLIGENT SSE APPLY TO THE ZERO-TRUST MODEL?

Now that we have covered the core principles of zero trust and the supporting security model where the intersection of data with other components is critical, we can map out how Netskope Intelligent SSE adds value.

First, implicit trust is removed as each session has transactional risk assessed for the relevant component risks. Netskope supports application risk profiles, user risk profiles, and device posture checks, plus the exchange of risk profiles with third-party security solutions, such as the Tanium endpoint management solution with thousands of posture checks in a near real-time assessment or CrowdStrike as an EDR/XDR solution in real time. Netskope integrates with leading IAM solutions for identity services and multi-factor authentication (MFA) with the ability to request step-up authentication for the transactional risk of a session. For example, a user wants to delete data in a company-managed application and their user risk score is medium, so the real-time adaptive policy can invoke a step-up authentication request before the data is deleted.

Second, least privilege providing the bare minimum of access goes well beyond static allow-and-block policy controls with adaptive access controls based on rich context and user input. For example, users can be coached in real time to select safer alternative applications from risky ones, asked to provide a justification for their higher risk choice, or alerted to the risky transaction with the option to cancel it. Least-privilege access is further enhanced by understanding over a hundred unique detailed activities for thousands of applications with Netskope. For example, Slack has 15 activity controls and Zoom has 10 activity controls in Netskope SSE compared to legacy solutions with just allow or block controls. Instance awareness for over 480 applications provides an understanding of the company instance versus a personal instance for the same application. For example, with legacy controls you allow Google Workspace and Google Drive without an understanding of instance. This enables users to exfiltrate data to their personal instance of Google Drive, whereas Netskope understands the difference and can provide real-time coaching or prevent exfiltration of sensitive data to personal instances. For applications that do not support instance awareness, identity mapping extends the use case for company and personal access.

Third, comprehensive security monitoring goes beyond detailed transaction event streaming with business intelligence analytics and visualizations to uncover the unknowns for zero-trust, and provide least-privilege policy refinement. Dashboards and charts for potential insiders, risky users, data movement between application instances, application risk profiles and trends, user behaviors, plus advanced threat and data protection reporting provide a closed loop back to adaptive access policies for refinement. The same rich context that enables least-privilege adaptive access controls is provided in advanced analytics for three, six, or 13 months for year-over-year analysis.

These are users with at least one event from the three risk indicator categories below

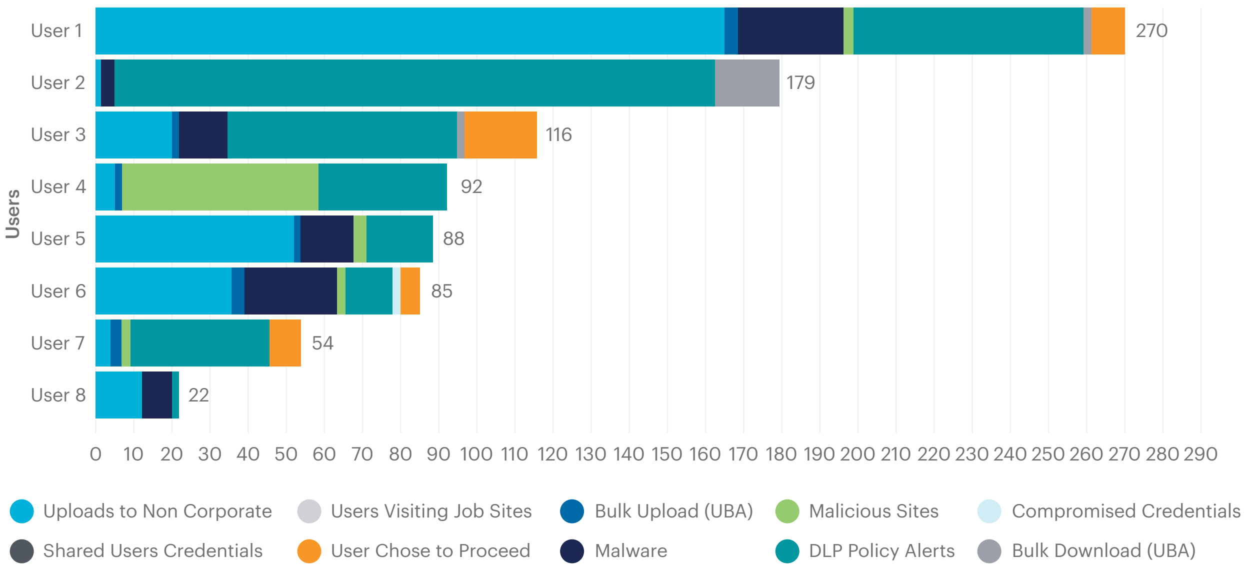

Figure 2: Top risky users as potential insider threats with three or more risk indicators.

Assessing the three core principles of zero trust shows the depth of Netskope Intelligent SSE beyond legacy perimeter defenses built around allow or block policy controls for ports, protocols, domains, URLs, and applications. Decoding applications inline to understand the rich context and content alongside transactional risk with the ability to alert and coach users to make real-time decisions for the session is the new set of requirements for zero-trust traffic inspection. Like zero trust itself, the focus shifts to use cases for business outcomes beyond the core capabilities of data and threat protection, acceptable use of the web and cloud, and TLS traffic inspection. In the diagram below are six popular use cases for SSE based on requests and popularity from customers.

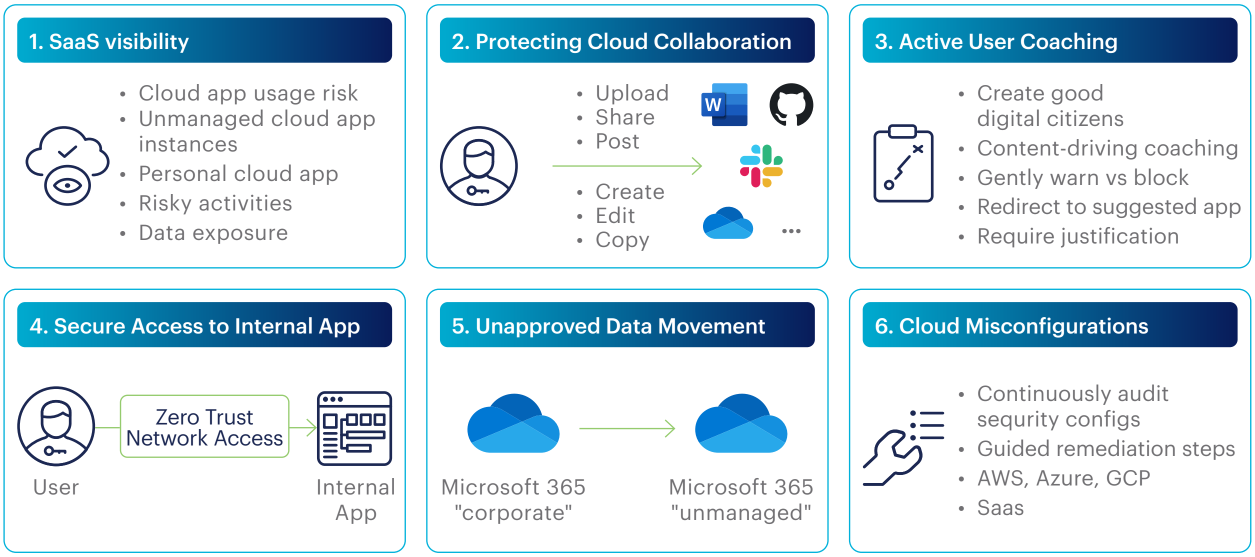
1. SaaS visibility
Cloud app usage risk
Unmanaged cloud app instances
Personal cloud app
Risky activities
Data exposure
2. Protecting Cloud Collaboration
Upload
Share
Post
Create
Edit
Copy
3. Active User Coaching
Create good digital citizens
Content-driving coaching
Gently warn vs block
Redirect to suggested app
Require justification
4. Secure Access to Internal App
Zero Trust Network Access
User
Internal App
5. Unapproved Data Movement
Microsoft 365 "corporate"
Microsoft 365 "unmanaged"
6. Cloud Misconfigurations
Continuously audit sequrity configs
Guided remediation steps
AWS, Azure, GCP
Saas

Figure 3: Six high-value use cases for Netskope Intelligent SSE.

Next is mapping Netskope Intelligent SSE solutions and capabilities to the zero-trust security model for the intersections of data with users, devices, applications and workloads, and networks. Since we reviewed the broader scope of these intersections previously in the paper, here is a summary listing of capabilities.

Users & Identities:

- Integration with leading IAM solutions for SSO and MFA capabilities
- Step-up authentication based on transactional risk for each session
- UEBA using 109+ detectors and 62+ ML models for users and peer groups
- User Confidence Index (UCI) risk scoring per user activities, events, alerts, etc.
- UCI score assessed for transactional risk for each session, plus workflow automation
- Instance awareness of users in hundreds of applications plus identity mapping of logins
- Real-time coaching to users with the ability to collect justifications

Devices & IoT:

- EDR agent posture checks in real time (e.g., CrowdStrike)
- Thousands of security posture checks for devices in near real time (e.g., Tanium)
- Discovery, classification, risk assessment, and control of IoT assets
- Device network location for company network versus remote working
- Single client for web and cloud access, plus ZTNA to private apps and resources
- ZTNA Next combines an SD-WAN client and a ZTNA client into one to replace VPNs

Applications & Workloads:

- CASB inline for thousands of applications and CASB API for managed applications
- Instance awareness (company versus personal) for hundreds of applications
- User identity mapping for non-instance applications to control access
- Data and threat protection for inline and API inspection of applications
- Risk profiles for 80,000+ applications (Cloud Confidence Index ratings)
- 100+ unique application activity controls for inline policy controls
- SSPM for configuration checks, compliance, and posture of SaaS
- CSPM for configuration checks, compliance, and posture of IaaS/PaaS
- Dedicated egress IP addresses to enhance SaaS security

Networks:

- ZTNA for private applications and resources, plus campus ZTNA and browser access
- Clientless browser ZTNA for contractors and third parties to apps and resources
- Advanced DLP for web, cloud, email, and endpoints
- Threat protection with AV, ML classifiers, sandboxing, and patient-zero protection
- TLS inspection of web and cloud traffic with a 50 ms SLA for round-trip time
- SWG for HTTP/S web traffic with 120+ web filtering and application categories
- Browser isolation for risky websites, unrated sites, NRDs, NODs, and parked domains
- Cloud firewall (CFW) of egress traffic with application firewall rules and DNS security
- IPS for web and non-web traffic including client traffic exploit protection of malicious websites
- Digital experience management (DEM) for user experience, device, SSE (time in Netskope), and network and application performance
- 70+ NewEdge regions across the globe, full compute, extensive peering, no surcharges (except China)
- Transaction event streaming, plus advanced analytics with graphic visualizations and dashboards

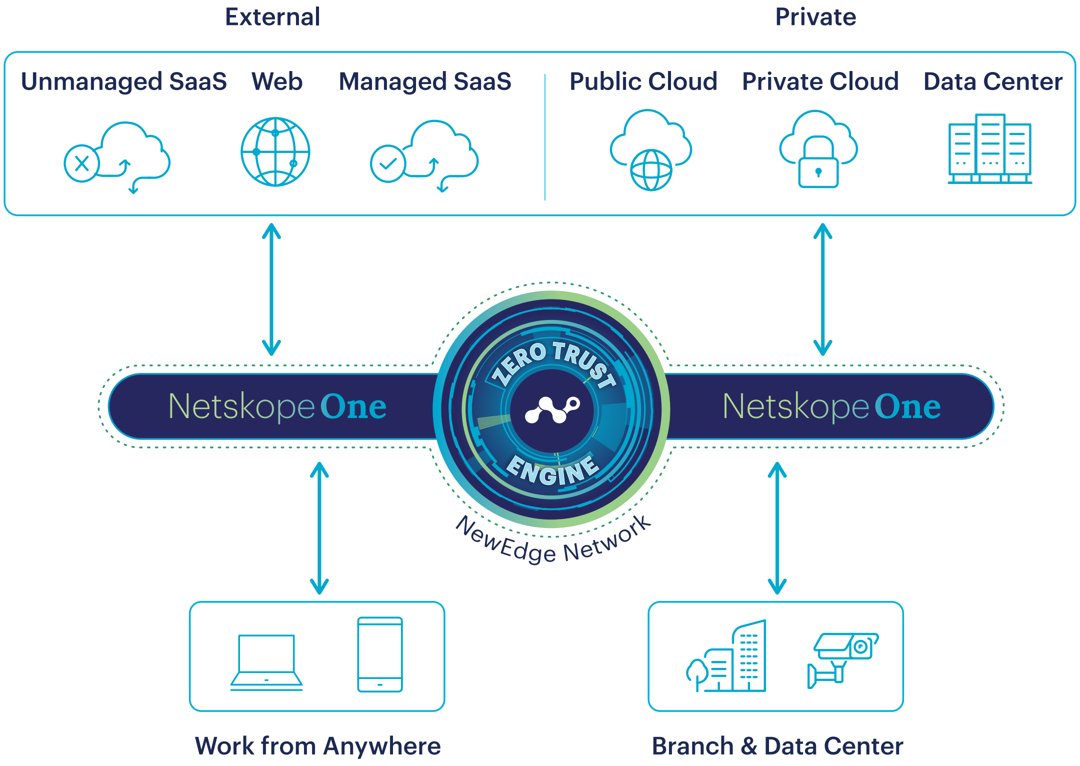

Figure 4: Netskope One architecture.

Netskope Intelligent SSE supports zero-trust principles and the zero-trust security model with one cloud platform, console, policy engine, client, and analytics. The single-pass architecture avoids service chains and the resulting latency with full parallel computation of all services on data planes to assess transactional risk for each session. At the core of the architecture is the Zero Trust Engine, a topic we will review in detail after integration with existing security infrastructure.

HOW DOES NETSKOPE INTELLIGENT SSE INTEGRATE WITH EXISTING SECURITY INFRASTRUCTURE?

For 1:1 or 1:M (Many) integrations, often published APIs solve the integration issue between security solutions given you have the IT resources. However, for M:M integrations and organizations short of IT staff for API projects, Netskope provides Cloud Exchange with four modules and 70+ integrations ready to use with a plug-in architecture. Cloud Exchange was designed to ease and scale integrations and is no charge to customers . Customers can deploy Cloud Exchange and its modules in Docker, Kubernetes, or OpenShift environments. Each Cloud Exchange module solves a different use case as follows:

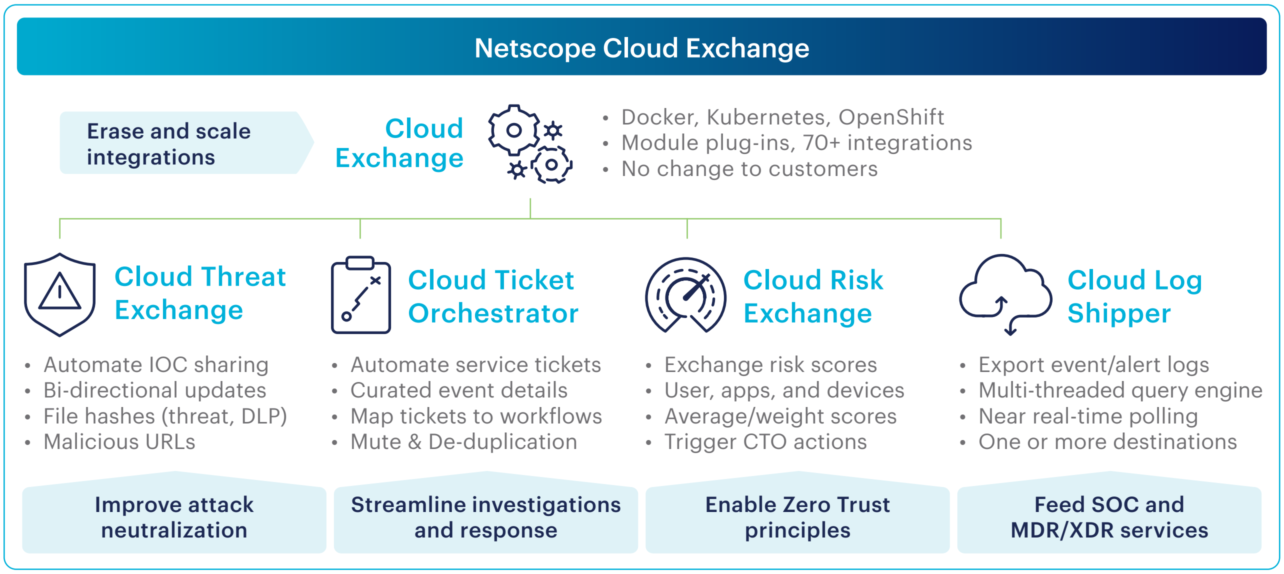

Netscope Cloud Exchange

Erase and scale integrations

Cloud Exchange

- Docker, Kubernetes, OpenShift
- Module plug-ins, 70+ integrations
- No change to customers

Cloud Threat Exchange

- Automate IOC sharing
- Bi-directional updates
- File hashes (threat, DLP)
- Malicious URLs

Improve attack neutralization

Cloud Ticket Orchestrator

- Automate service tickets
- Curated event details
- Map tickets to workflows
- Mute & De-duplication

Streamline investigations and response

Cloud Risk Exchange

- Exchange risk scores
- User, apps, and devices
- Average/weight scores
- Trigger CTO actions

Enable Zero Trust principles

Cloud Log Shipper

- Export event/alert logs
- Multi-threaded query engine
- Near real-time polling
- One or more destinations

Feed SOC and MDR/XDR services

Cloud Threat Exchange (CTE) improves attack neutralization by automating bi-directional IOC updates between customer defenses, including file hashes and malicious URLs. This enables a M:M integration between SSE, EDR, SEG, XDR, SIEM, SOAR, and threat intelligence sources for customer environments. CTE has also been used for DLP file hash sharing as another use case.

Cloud Ticket Orchestrator (CTO) automates and maps service tickets to workflows with curated event details to streamline investigations and response. Popular IT service management solutions like ServiceNow, PagerDuty, and Jira have ready-to-use plug-ins, plus collaboration solutions including Slack and Microsoft Teams.

Cloud Risk Exchange (CRE) enables the exchange of risk scores for users, apps, and devices between security solutions with the ability to average and weight scoring and trigger CTO actions. This enables customers to share application risk scores with security assessment solutions, exchange EDR device risk scores, IAM identity risk scores, and more. SSE adaptive access policy controls utilize risk scores in real time for desired transactions to support zero-trust principles.

Cloud Log Shipper (CLS) automates event and log export to third-party solutions using a multi-threaded query engine with near real-time polling. CLS can feed 1:M destinations including XDR, MDR, SIEM, and SOC operations. Netskope also provides transaction event streaming directly from the SSE platform to destinations as another option beyond CLS.

Beyond Cloud Exchange and its modules, Netskope provides REST APIs for direct integration, scripting, and automation.

WHAT IS THE NETSKOPE ZERO TRUST ENGINE (ZTE)?

At the heart of Netskope Intelligent SSE is the Zero Trust Engine (ZT Engine) for inline real-time adaptive access controls for user traffic to web, SaaS, unmanaged SaaS, cloud service providers for IaaS and PaaS, and public-facing private apps. To support zero-trust principles and the zero-trust security model, ZT Engine assesses many variables at the time of the business transaction, provides real-time coaching to users, collects justifications, and logs events with rich details for continuous monitoring to find unknown risks to refine policies in a closed loop. Figure 5 provides a more detailed look at ZT Engine and its variables for adaptive access controls, coaching users, and policy refinement.

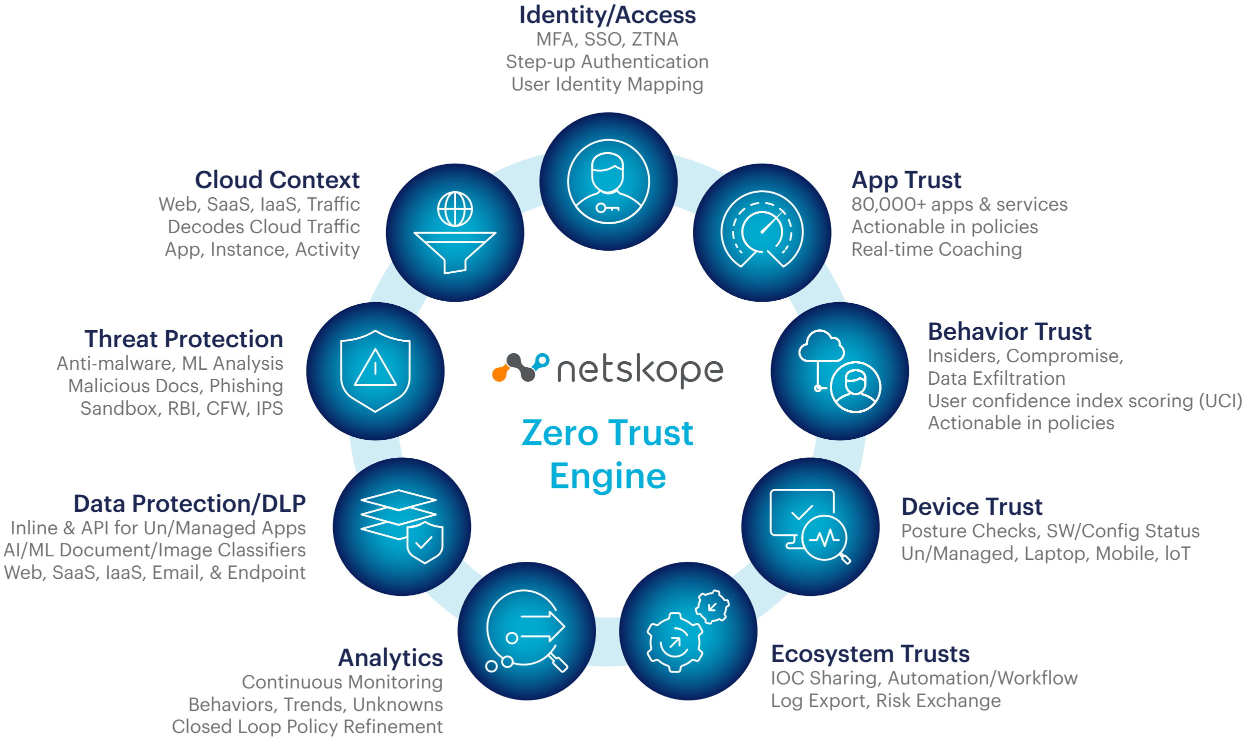

Figure 5: Netskope Zero Trust Engine providing risk-based context for zero-trust policies.

Starting clockwise at the top with Identity/Access is the integration with IAM and IdP solutions as a true source for an identity to ZT Engine. Multi-factor authentication (MFA) and single sign-on (SSO) enhance security access as well as replacing VPNs with ZTNA as discussed earlier in the paper. ZT Engine also tracks identity mapping per user to extend the use case of instance awareness when applications do not provide instances. Next are variables for application, behavior, and device trust, all with defined risk scoring and security checks into ZT Engine. Does the application have a low score for its Cloud Confidence Index (CCI) and should the user be coached to a safer alternative? Has past behavior of the user or entity resulted in a low User Confidence Index (UCI) score? Should the device type and posture be trusted? Plus, the security ecosystem integration of Cloud Exchange modules provide input to ZT Engine with IOC sharing, risk score exchanges, and automation of workflows.

Moving back up to the top of the ZT Engine diagram and going counterclockwise from Cloud Context, we hit upon a very important advancement for zero-trust principles including least-privilege access in adaptive access policy controls. Cloud context is derived from decoding inline application transaction requests to learn the application context, content, activity, and instance. This enables ZT Engine to understand data movement between company and personal instances of an application often associated with employees gathering useful information before departing an organization. Also, the requested activity such as deleting numerous files with remote access from a mobile device in a company-managed application where a step-up authentication request can be invoked by ZT Engine to verify the identity a second time. Cloud context is the missing piece for zero-trust principles including least privilege with legacy security defenses including secure web gateways (SWGs) built around allow/deny URL filtering and file controls, or next-generation firewalls (NGFWs) with allow/deny controls by app, user, and content. If you deny app access with legacy defenses, you impede digital transformation progress, and an open allow enables data exfiltration and data theft to personal instances to go unmonitored.

Next are threat and data protection defenses into ZT Engine where increasingly AI/ML advancements are improving efficacy and time to value. For data protection and DLP, ready-to-use AI/ML classifiers for 27 documents and images can detect over 20 types of source code, resumes, screen captures, and whiteboards, plus legal, financial, and medical documents. Data protection can also put guardrails on data movement to high-risk applications and to personal instances, or unmanaged applications. One of the key values of real-time coaching with data protection is advising users on risky data activity where more than 95% of users will avoid the risk and cancel the transaction, and for the remaining 5% a justification can be collected to learn why and how to refine policies for these use cases. Security teams can avoid an overall blocking policy, educate users, learn about exceptions, and reduce help desk tickets often associated with blocking policies.

For threat protection, new AI/ML classifiers for PE files, PDF files, Microsoft Office files, and for phishing attacks are improving efficacy beyond what signatures, sandboxing, egress firewalls, intrusion prevention systems, and remote browser isolation provide. AI/ML is also utilized for dynamic webpage classification, detecting domain generation algorithms (DGA), and enables ZT Engine to make policy decisions to protect users against unknown zero-day threats.

The final part of ZT Engine is key to the last principle of zero trust to continuously monitor and refine least-privilege access. Most legacy security solutions are content to stream and off-load event logs to third-party solutions such as a SIEM or data lake. However, key to ZT Engine is cloud context, understanding data sensitivity, app risk, instance, and activity. This enables advanced analytics driven by high-end business intelligence visualizations such as the Sankey chart below showing data movement for a user, their identities, applications, instances, and activity.

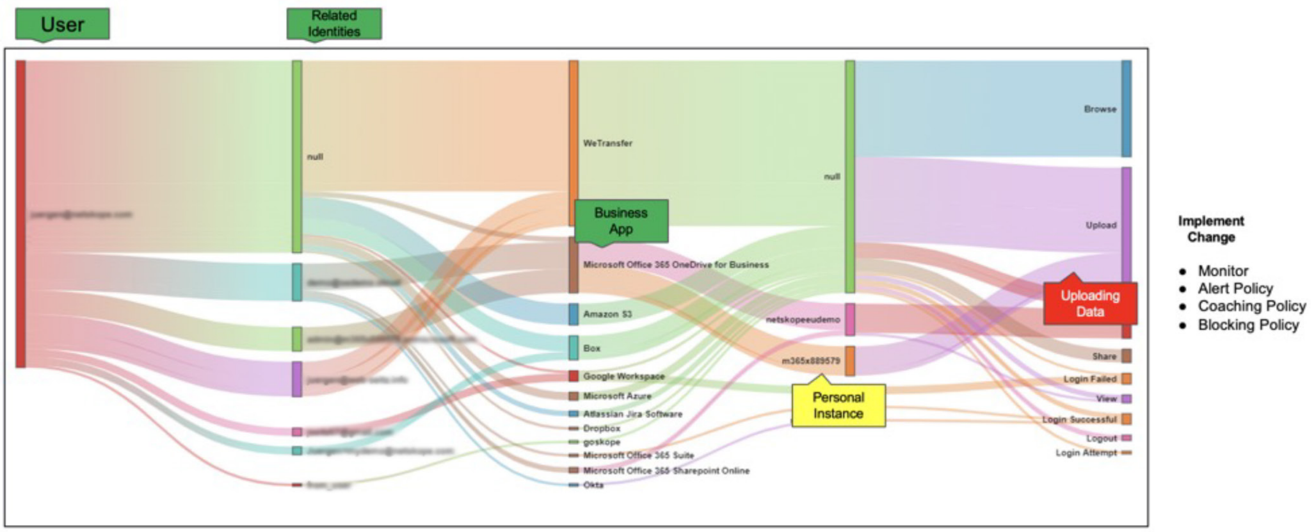

Figure 6: Netskope Advanced Analytics providing visibility of unknown data exfiltration to personal storage.

This addresses one of the key use cases of detecting unknown and unapproved data movement, as the flow of data between zero-trust elements is core to the zero-trust security model. While deployed policies can address least-privilege access for known risks, the continuous monitoring through advanced analytics enables policy refinement for unknown risks in a closed loop that raw logs and queries are unlikely to expose. This puts the value of rich context collected by ZT Engine into high-value visualizations, dashboards, and drill-down details for a more efficient and effective SSE solution supporting zero-trust principles. To see a short demo of Advanced Analytics, visit this webpage.

HOW DOES ZT ENGINE INVOKE ADAPTIVE ACCESS CONTROLS AND WHAT ARE SOME COMMON USE CASES USING ZT ENGINE?

Generative AI has had a firestorm of attention, and security teams are reacting quickly as employees, contractors, and business partners leverage the value and economies of scale from this new application. However, there are risks for what data gets provided to generative AI applications like ChatGPT, Gemini, and CoPilot. The popularity also provides an opportunity for cyberattacks to phish, lure, and collect data from oblivious users. Generative AI is a good use case to illustrate the value of ZT Engine to provide safe access, protect data, and prevent threats related to its popularity.

Blocking access to generative AI apps via specific domains or a URL category prevents progress and advancements from using it across many functional areas of an organization, including developing code, creating marketing content, and improving customer support. However, to openly allow access to generative AI apps puts sensitive data at risk from uploads for analysis, and not every AI application or front-end hosted website is low risk and well managed. So, a gray area develops of allowing access with guardrails for data protection, understanding company instances, activities, and when to provide real-time coaching to users. ZT Engine has app connectors for generative AI applications to decode the application transactions inline to provide adaptive access policy controls, protect data, and coach users in real time.

For example, a user with a company identity who is accessing ChatGPT can be provided an alert on best practices for using generative AI and to avoid submitting sensitive company data. As the user interacts with ChatGPT, data exposed to the application can be inspected for data sensitivity using standard, advanced, and AI/ML classifiers for data loss protection. The company may have also invested in its own private version of the application, and this company instance can have different ZT Engine policy controls than the public version. Monitoring overall use of generative AI applications can be provided by a URL category plus the rich context collected by ZT Engine in advanced analytics for specific application transactions, including cloud context for the app, instance, activity, data sensitivity, collected justifications for use, and any policy alerts. As noted earlier, ZT Engine can assess the app risk, device posture, and user behavior risk score, plus any threat or data protection defense inputs to make an adaptive access policy decision in real time.

So, it is no surprise that data flowing into generative AI applications is at the center of attention for risks to organizations. This use case and many others show why data is at the center of the zero-trust security model as it interacts with users, devices, apps, and networks. This makes context and content the new perimeter for users and apps working from any location where hybrid work is the new normal. One of the main goals of implementing zero trust is to make security a business enabler so your organization is more trusted, projects can leverage digital transformation safely, and data is protected in use, in motion, and at rest.

WHAT ARE THE OUTCOMES AND BENEFITS OF ADOPTING SSE AND ZERO TRUST?

Netskope surveyed hundreds of its customers to learn the outcomes and benefits for adopting SSE solutions and zero-trust principles. Three key outcomes surfaced from the survey, with the first being an 85% reduction in security risk through critical asset protection, stability, resiliency, and making users better digital citizens. Next was a 51% reduction in total cost of operations by retiring appliances and freeing up full-time employees, reducing dedicated network links, improving operational efficiency with consolidation, and optimizing cloud spend. And as expected by moving defenses closer to users and adopting zero-trust principles, there was a 19% improvement in business agility with an improved user experience, time to value, and speed to market, plus making data-driven decisions.

Netscope Business Value Benefits

The following benefits are based on actual netscope customers data

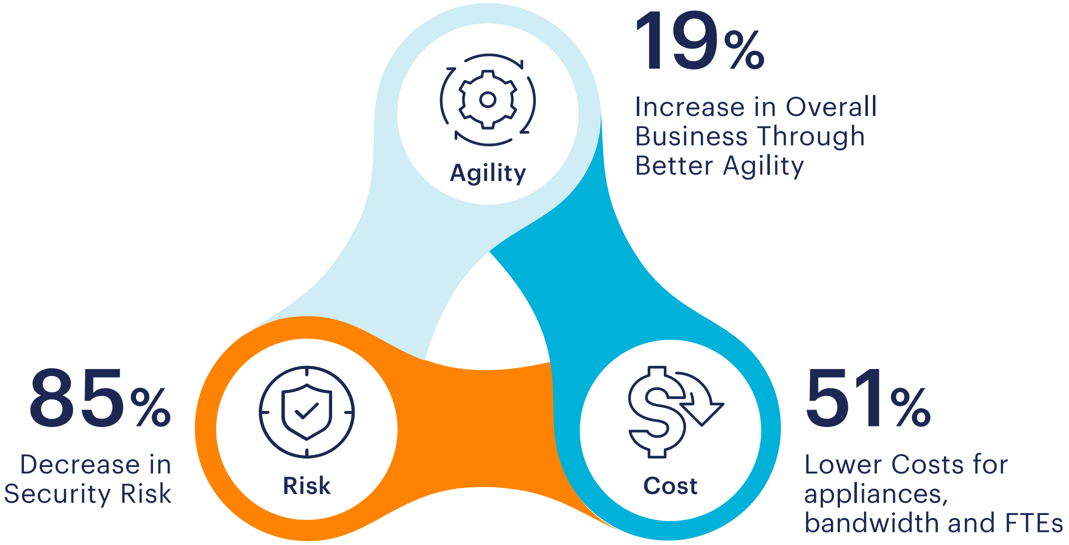

19%
Increase in Overall Business Through Better Agility

85%
Decrease in Security Risk

51%
Lower Costs for appliances, bandwidth and FTEs

As noted earlier in the paper, companies that have adopted zero-trust principles often exceed compliance regulations and find it easier to support new business and operational models. Security becomes a business initiative, not a set of technologies or an expense. Zero-trust principles start with business initiative strategies, objectives, and tactics to enable agility, build trust and confidence, and differentiate solutions and services from competitors. The outcomes and benefits are hard to ignore, even more so as digital transformation divides the leaders from the pack to survive.

WHAT GETS REPLACED OR RETIRED WITH A ZERO-TRUST SSE TRANSFORMATION?

The transformation impacts four areas when you intersect SSE with zero-trust principles for networking, security, applications, and data. You can start in one or more of the areas; however, a key to success is having cloud context and content visibility and control as noted earlier for the Zero Trust Engine. Legacy allow and deny controls will impede your success and use cases in the four transformation areas for both SSE and zero-trust principles and impact your SASE architecture.

First is networking, as the impact is substantial for improving user experience, reducing costs, and providing least-privilege access. Your goals in this transformation stage are supporting a hybrid work-from-anywhere environment, a great user experience, federating MFA to a growing base of SaaS applications, and zero-trust network access (ZTNA) to private apps and resources removing the open service liabilities of VPNs and lateral movement. The monetization of using remote access credentials by ransomware alone is a reminder why zero-trust principles are required. The transformation also includes reducing spend on MPLS and dedicated WAN links, moving to a direct-to-cloud access design with SSE cloud platforms close to any user location, and adopting SD-WAN, including in managed endpoints with ZTNA.

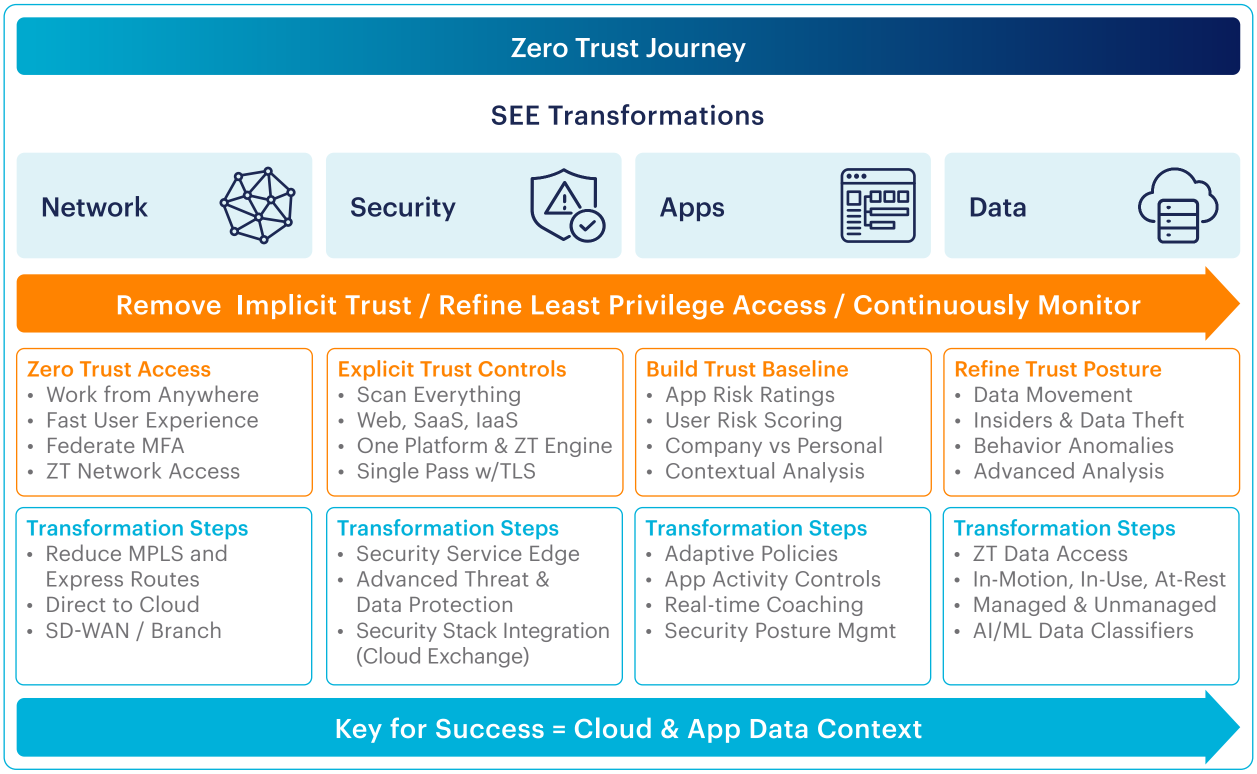

Zero Trust Journey

SEE Transformations

Network

Security

Apps

Data

Remove Implicit Trust / Refine Least Privilege Access / Continuously Monitor

Zero Trust Access
- Work from Anywhere
- Fast User Experience
- Federate MFA
- ZT Network Access

Explicit Trust Controls
- Scan Everything
- Web, SaaS, IaaS
- One Platform & ZT Engine
- Single Pass w/TLS

Build Trust Baseline
- App Risk Ratings
- User Risk Scoring
- Company vs Personal
- Contextual Analysis

Refine Trust Posture
- Data Movement
- Insiders & Data Theft
- Behavior Anomalies
- Advanced Analysis

Transformation Steps
- Reduce MPLS and Express Routes
- Direct to Cloud
- SD-WAN / Branch

Transformation Steps
- Security Service Edge
- Advanced Threat & Data Protection
- Security Stack Integration (Cloud Exchange)

Transformation Steps
- Adaptive Policies
- App Activity Controls
- Real-time Coaching
- Security Posture Mgmt

Transformation Steps
- ZT Data Access
- In-Motion, In-Use, At-Rest
- Managed & Unmanaged
- AI/ML Data Classifiers

Key for Success = Cloud & App Data Context

Second is security transformation by moving away from legacy security appliances to an SSE cloud platform with a single-pass TLS inspection path with no tradeoffs between performance and security. Scan everything possible with no bypasses around office productivity suites, some that lead the market for cloud malware delivery while others are the leading location for data exfiltration and theft as employees depart. Using one consolidated SSE platform, analyze web, SaaS, unmanaged SaaS, cloud service provider access, and public-facing private apps recognizing more than half of malware and threats are cloud-delivered today. Leverage advanced threat protection and data protection increasingly using AI/ML-based defenses and classifiers in real time for inline and out of band for API inspection. Security stack integration is also important for sharing threat intel and IOCs between your security defenses, exporting logs and events, exchanging risk scores for apps, users, and devices, plus automating workflows into IT service management and collaboration solutions.

Third is application transformation where cloud context opens new capabilities for adaptive access controls and real-time coaching often not seen with legacy allow and deny security solutions. Application risk scores and cloud risk assessments can align users with safer applications, company-managed applications, and advise users to avoid high-risk applications unless they provide a justification. This transformation stage also adds behavior anomalies (UEBA) from cloud context to understand past behavior to apply a risk score used in adaptive access controls per entity or user. And as noted for the Zero Trust Engine is contextual least-privilege access based on the app, app risk, app instance (company versus personal), activity, user risk, data sensitivity, and other contextual variables about the device and location. This expands use cases to adaptive access policy controls, real-time coaching to users, and applying application activity controls including invoking step-up authentications. Part of the application transformation stage also includes security posture management for applications and cloud services where the first starts out mostly wide open and the second is secure until it is not.

Fourth is data transformation where you gain a better understanding of data movement. Pre-pandemic it was likely your users, applications, and data were on your network using your data centers with a small amount of remote access and SaaS adoption. Post-pandemic and we increasingly see hybrid work environments, an increase in SaaS adoption year over year, and data moving out into the cloud with these users and applications. Data needs guardrails between applications, instances, users, and activities that SSE can deliver with cloud context for adaptive access controls and real-time coaching before you invoke formal DLP. Cloud context combined with UEBA machine learning also helps detect insiders, data theft, and unknown data exfiltration as shown above in the advanced analytics Sankey (flow visualization) chart. Data is the key element that connects and flows with all the other parts (users, devices, networks, applications/workloads) of the zero-trust information security model.

WHAT IS THE FEEDBACK FROM NETSKOPE CUSTOMERS?

One source for reviews directly from customers is Gartner Peer Insights, where SSE solution reviews for Netskope can be viewed at https://www.gartner.com/reviews/market/security-service-edge/vendor/netskope. To highlight feedback from a CISO of a large IT services company, here are the direct quotes from the review:

> “Netskope security cloud provides us visibility into user activity on the web that was previously impossible with legacy technologies. With this, we are able to provide secure business enablement — allowing a global approach to web/cloud security that's fast and effective anytime and anywhere. With this additional context, we can now do things like:
>
> - Allow seamless work with third parties and their cloud apps
> - Add flexibility to our employees for mixed personal/work usage
> - Allow global remote working
>
> All without the added risk this used to bring.
>
> Netskope's granular analysis and real time protection allows us to have our cake and eat it too — better security and better user experience simultaneously.”

In the same review, the CISO shared some lessons learned that are worth a review, as they are consistent with other customers' feedback.

- “Granular visibility into cloud apps and activities rather than just category and URLs.
- Seamless global routing to nearest POP for ensuring a fast-browsing experience no matter which office or country employees visit.
- Understanding of instance awareness. And being able to control data flows based on that context.
- Single management console for all employees, cloud apps and web activity globally.
- Out of the box risk analysis of tens of thousands of cloud apps being used by employees and being able to make policy decisions in real time based on risk.
- A dedicated TSM team or individual is given to our organisation, allowing us a single point of contact and technical resource included for free — a very personal relationship can be formed to save time on complex issues and ensuring we take the product to its maximum value.
- Conditional access into Private Apps, taking into account device context, significantly bolstering security of private apps.
- Netskope Private Access can be used even internally as a tunnel into Private Apps, using user to app style segmentation, instead of standing source IP ACLs that create risk. This allows a Zero Trust approach and bolsters network segmentation efforts significantly.”

The review also included the reasons driving the change to use Netskope.

- *Improve compliance & risk management.*
- *Improve business process agility.*
- *Improve business process outcomes.*
- *Enhance decision making.*
- *Drive innovation.*
- *Create internal/operational efficiencies.*

There are challenges in the reviews and one of them is learning how to leverage the context of the Zero Trust Engine in adaptive access controls to support the use cases desired. Transforming from a mindset of allow-and-deny controls around network access and web filtering to understanding how applications can be safely utilized and how to protect data flowing between apps, users, and devices is a shift in thinking. Generative AI applications like ChatGPT were noted earlier as a prime example to drive innovation while protecting company sensitive and confidential data. The world of my network, my device, my applications, and my data is fading to where the only part you may own is the data as a remote user with a personal device on a public network accesses a managed SaaS application. What many of the reviews will highlight is the granular visibility and context from Netskope SSE that makes securing hybrid work environments easier with a better user experience while advancing to a zero-trust security model.

Following the same theme of granular context and visibility, this CISO of a large enterprise retail company noted how Netskope is helping their organization.

> "Obtaining visibility of analytics not previously available with our prior solution and improving our zero-trust security posture."

The Zero Trust Engine provides real-time adaptive access controls based on rich cloud context, and real-time coaching to users as consistent themes in customer feedback noted by this information security leader of a financial services firm using Netskope SSE.

> "The fine-grained access controls by creating real-time policies which are tied to our IDP have enabled us significantly better in terms of improving our zero-trust network."

The intersection of SSE and the zero-trust security model is a strategic and important decision for every company. Zero-trust principles require rich context for least-privilege access and for continuous monitoring to find unknown risks to loop back and improve least-privilege access. If you only address the first principle of removing implicit trust, you are thinking tactically. Here is a quote from a Fortune 500 financial services firm architect who plans the future success of their company concerning this decision.

> "Netskope allows for the enablement of our Zero Trust strategy."

WHAT ARE THE INDUSTRY ANALYSTS’ PERSPECTIVES?

If you are reviewing the latest analyst reports year over year for SSE and wonder why one vendor moves forward and the other vendors backwards, you are witnessing the transformation to a zero-trust security model. A legacy perimeter mindset built around firewalls, VPNs, and secure web gateways is unable to decode application traffic inline and will miss employees departing with sensitive data in personal cloud storage applications or safely enable generative AI applications while protecting sensitive company data. Shifting to an application and data-centric security model using cloud-edge security for any user, location, or device involves retiring old beliefs and recognizing the changes upon us.

We encourage readers to view the latest analyst reports about SSE for market share and vision, plus critical capabilities and leading use cases. Netskope rated first for two use cases and second and third for the remaining two use cases, being the highest ranked SSE leader in the reports. You can download the analyst reports from https://www.netskope.com/resources/analyst-reports.

ABOUT THE AUTHOR

Tom Clare has been in product strategy, management, and marketing of security solutions for over 25 years covering secure web gateways, firewalls, intrusion prevention systems, antivirus/antimalware solutions, deception, UEBA, network detection and response, MDR, plus encryption solutions. He was the managing editor and contributor for the book on Borderless Behavior Networks, now in its second edition. Presently he is a Senior Director at Netskope focused on SSE solutions and is certified on Applying Zero Trust.

Interested in learning more?

Request a demo

Netskope, a global SASE leader, helps organizations apply zero trust principles and AI/ML innovations to protect data and defend against cyber threats. Fast and easy to use, the Netskope One platform and its patented Zero Trust Engine provide optimized access and real-time security for people, devices, and data anywhere they go. Thousands of customers trust Netskope and its powerful NewEdge network to reduce risk and gain unrivaled visibility into any cloud, web, and private application activity—providing security and accelerating performance without compromise. Learn more at netskope.com.