netskope

WHITE PAPER

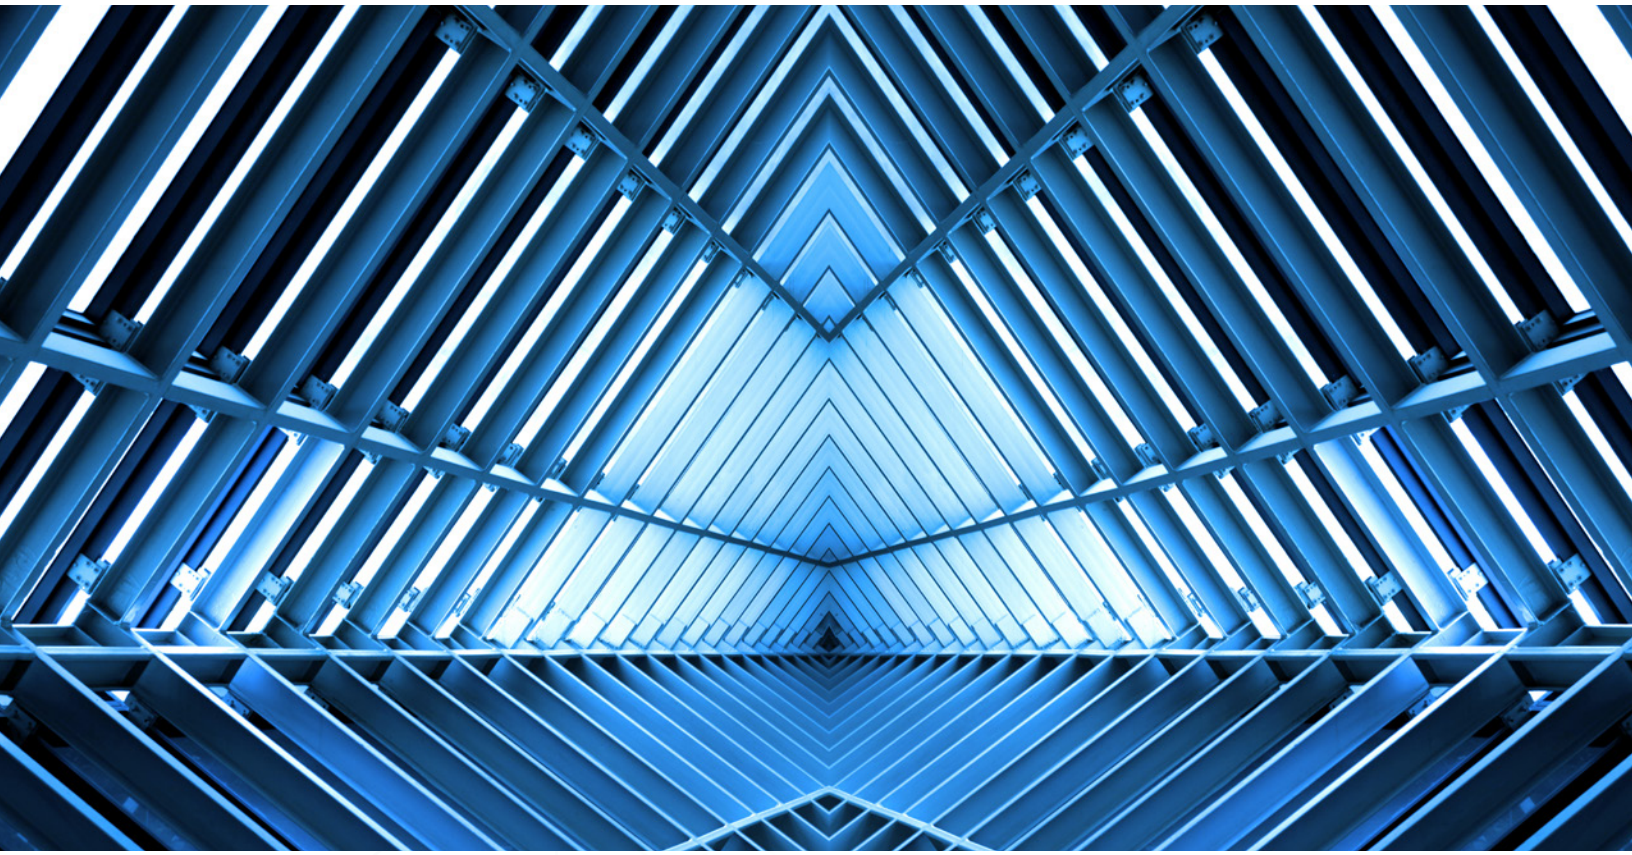

Netskope Reference Architecture for Zero Trust

A Technical Guide to Zero Trust Principles, Pillars, and Capabilities for Federal Government Agencies

PURPOSE

INTRODUCTION

Zero trust is a concept of applying the least amount of trust possible to allow an entity the ability to access the resource it requires to complete the assigned or needed task. Zero trust works on the principle there is no implicit trust or baseline access to a resource, removing the concept of trusted by location as seen with the classic perimeter defense system.

Netskope's Zero Trust Architecture (ZTA) models standards outlined in the Department of Defense Zero Trust Reference Architecture, NIST publication 800-207, and the National Security Agency's Cybersecurity Information Sheet: Embracing a Zero Trust Security Model.

PURPOSE

Netskope defines Zero Trust simply as a security model based on the premise that no one should be blindly trusted inside the network or allowed access to resources, applications, or data until they have been validated as a legitimate user/entity with a legitimate need to access the information/data. Netskope employs Zero Trust as an operation concept to the services provided as part of the secure access service edge (SASE) architecture built on the Netskope NewEdge network. This data-centric design allows for maximum interoperability across all applications while enforcing trust based on context, at the network, application, and user level, irrespective of a user's location or access method.

OPERATIONAL CONCEPT

Netskope implements its ZTA in a layered approach divided into three functional groups: Identity and Access Method, Resource, and Control. This design provides a functional grouping of activity and helps align integrations with technology partners and functions as a reference to understand flow of a request.

A request for a resource is a linear progression through the operational model where identity and access method are first understood and then policy enforcement takes place at one or more levels as the user/entity is gaining access to the resources. The resources are defined as three types: Resource, Application, and Data, each having an associated policy enforcement point (PEP) to allow the appropriate level of access to the resource. Allowing the user/entity access to a resource is first completed by challenging the user/entity to authentication and understanding the method of the access request and posture of the device for access. Netskope then looks to understand the user/entity level of authorization; this is based on both static policy application as well as dynamic understanding of behavior. Control of policy, learning, reporting, and integrations through the Control Plan are set up to allow for a dynamic feedback loop that is adaptable and modifies controls as needed based on behavior or outside input.

Full or limited access to the resource is controlled through the PEP. PEPs are capable of doing more than a simple allow/deny: Based on their functional level and understanding of the resource they are securing, they can function up to the Application layer and beyond with cloud-based applications and data.

The below diagram illustrates the data flow and PEPs of the architecture:

Netskope Zero Trust Capabilities Mapping

Native to Netskope's Cloud Security Platform providing ubiquitous policy control from a single UI

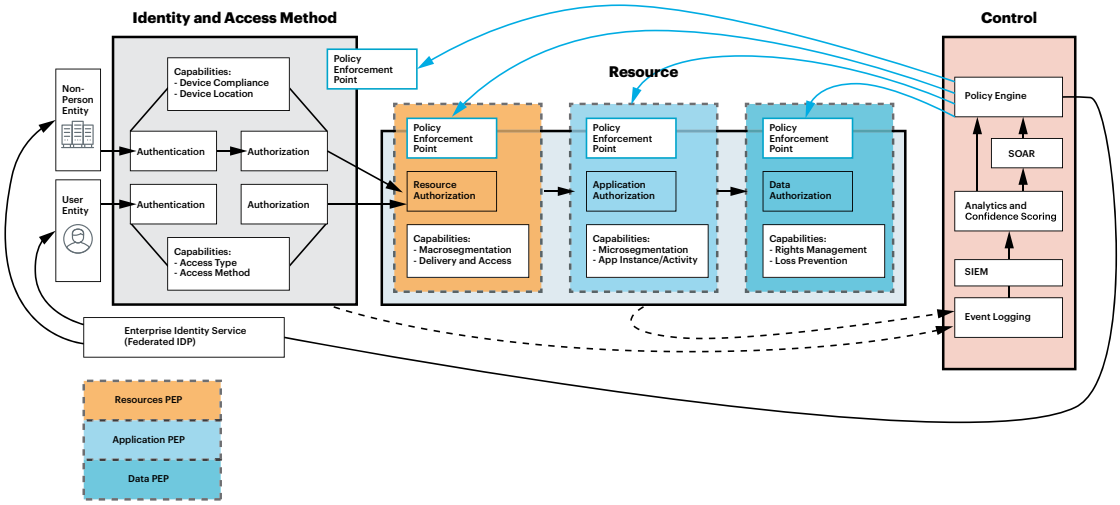

Identity and Access

The Identity and Access functional group is responsible for two major components that are enforced within this group and can be forwarded as part of the communication with the upstream Resource PEPs. The first responsibility is to authenticate the user/entity. This can be completed in concert with authenticating the device, providing a dual authentication, or it can be completed for just the user/entity in cases where the request is coming from an agentless device. The second core function of this group is to understand the authorization level of the user/entity.

Authentication

Authentication evaluates the identity of the user/entity (or non-person entity as per DOD use cases) and the device as access is requested to a resource, application, or service. Authentication is completed in a dual-factor mode, first by validating the device, next the user/entity. Device authentication is completed by validating the enrollment of the device with the Netskope platform. This is completed with a private public key infrastructure (PKI) enrollment during its initial setup as a managed device. User/entity challenges can be completed in several ways, most often via PKI or Security Assertion Markup Language (SAML) with integration to an Enterprise Identity Service. Authentication for unmanaged devices is also possible and functions as part of the Application request, most often through a web browser.

In the case of unmanaged devices authentication is only completed for the user/entity.

Authentication is tagged to the session as it flows to the Resource Group and it is evaluated at each upstream PEP to be used as part of policy evaluation for access.

Capabilities of Authentication include:

- Authentication of the device and verification of enrollment with Netskope
- Authentication of the user/entity with multi-factor authentication (MFA)
- Step-up authentication of the user/entity as part of the PEP

Authorization

Authorization evaluates the device posture and confirms the user/entity and their permissions to understand what level of access should be given. Integration often occurs with an Identity Provider. Authorization in the Netskope platform happens in this functional group as well as in the Resource functional group PEP; here Netskope is validating that the user has the access to initiate a connection to the Resource. Federated Identity services and Netskope policy, both static and dynamic, are responsible for setting the authorization level. Authorization also allows for understanding of the "state" of the request, defined below in capabilities. The state is passed to upstream PEPs to allow more granular enforcement of policy.

Capabilities of Authorization include:

- Determine if the device is managed/unmanaged
- Determine if security baseline requirements are in place
- Determine the Access Type
- Determine the Access Method
- Determine the Group Authorization—which group(s) the user and device have affinity to
- Dynamic authorization decisions based on user and device risk scoring

Resource

The Resource functional group is the logical grouping of the Network/Resource, Application, and Data resources along with their associated PEPs. Netskope's ZTA provides application of granular control at each of these resources.

Network/Resource Authorization

Network/Resource Authorization allows the Netskope Security Cloud to evaluate the request for access to understand access and allow PEP controls for:

1). Network Level Resources. These are resources or applications that are published to the user and not made available on the public internet. Access to this type of resource or application is determined by a network access control limiting the user/entity to only that resource restricting lateral movement.

2). Web Resources. Web locations/URLs are seen as resources and are available for policy application and forwarding to Browser Isolation as required based on the level of trust.

3). Software as a Service (SaaS) and Infrastructure as a Service (IaaS) Resources. These are public or private SaaS or IaaS services. Netskope has over 41,000 applications defined and the availability to create custom application connectors to support internal or private applications.

Capabilities of Resource Authorization include:

- Custom Content Categories
- Network Level Access (IP address)
- Microsegmentation
- Security and Activity Controls: Allow/Deny/Coach
- Remote Browser Isolation (RBI) targeted for uncategorized and security risk websites

Application Authorization

Application Authorization allows the Netskope Security Cloud to evaluate a request for an application resource (via Layer 7 protocol or application programming interface [API]). User entitlement is validated for access to private applications supporting certificate pinned applications and all TCP/UDP ports. This allows for application-level access without network access that can potentially lead to lateral movement. Each access permission is processed as a single flow providing additional visibility.

Capabilities include:

- Standard and custom categorization with Application Instance controls. Application instance is defined as the ability to understand the SaaS or IaaS to a level of abstraction where you are able to determine the unique tenant or instance the user is logging into. An example is determining if a user/entity is logging into their government account of Microsoft 365 or their personal account of Microsoft 365.
- Security and Activity controls:
    - Macrosegmentation
    - Allow/Deny/Coach/Notify
    - Multi-factor Authentication/Dual-factor Authentication, step-up authentication. and periodic re-authentication
    - Full API Controls; decode 70+ unique API controls such as login successful, login failed, logout, Create, Delete, Download, Edit, Follow, Post, Preview, Share, Upload, View (this is a subset depending on application)
    - All outbound TCP/UDP/ICMP traffic including DNS requests
    - Data context (DLP) and threat policy enforcement.
    - 5 tuple-based policy with outbound L3/L4 filtering for TCP, UDP, and ICMP with logging
    - Ability to create application definition, network event logging for non-web and web traffic
    - FQDN-based destination definition, FTP ALG, and Office 365 bypass
    - User identity with GRE steering
    - Inbound browser IPS/IDS
    - RBI targeted for uncategorized and security risk websites

Data Authorization

Data authorization allows the Netskope platform to provide validation of the data, at the application level or on the object before the user/entity is authorized access. Foundationally this is completed with data loss prevention (DLP) capabilities that allow data to be understood before access is given. Integrations with data tagging technologies allow application of policy enforcement of previously tagged data. Actions for application of data authorization PEP include granular enforcement based on application action, not only allow/deny/view only. Policy can be applied on the application instance or the application activity.

Capabilities include:

- Access permission and data context usage
- Dynamic risk evaluation and enforcement (time, context, user and entity behavior analytics [UEBA], device, location, group membership, device)
- Real-time remediation
- Includes 3,000+ data identifiers for 1,400+ file types, plus custom regex, patterns, and dictionaries
- File fingerprinting with degree of similarity and exact data matching
- Artificial intelligence (AI)/machine learning (ML) classification for documents (patents, M&A documents, tax forms, source code), plus images (desktop screenshots, passports, IDs)
- Incident management and remediation

Control

The control functional group is facilitated by native and REST API-based integration with enterprise infrastructure to provide visibility, enforcement points, and remediation for events and incidents supporting real-time, multi-mode inline steering for SaaS, IaaS, Platform as a Service (PaaS), web, and private applications.

The Netskope Security Cloud also provides an open platform with published APIs and a broad partner ecosystem landscape for integration. Security orchestration, automation, and response (SOAR) and security automation and orchestration (SAO) solutions can leverage the context from SaaS, PaaS, IaaS, and web transactions. This is used to ingest, enrich, and create response actions through playbooks. Integration is provided directly via REST API, JSON, CSV, or Key Value pairs format or native app via API tokens to pull event data (Events, Alerts, Anomalies, Violations, Threats, User Risk Score). Netskope Cloud Exchange (CE) is a near real-time threat ingestion, curation, and sharing system that enables up-to-the-minute intelligence feeds exchange with security infrastructure such as endpoints, firewalls, secure web gateway (SWG), and cloud access security broker (CASB). This also provides an integration platform for feeds into security information and event management (SIEM), API-based tools, or SOAR systems with a dashboard highlighting frequency original indicators of compromise (IOCs) hits, event correlation, scope of attack surface, and the ability to timeout threat indicators and observables due to staleness.

Capabilities include:

- Policy Engine
  - Policy enforcement
  - Share threat intelligence
  - STIX/TAXII capable
  - Adhere to API Rate Limits
  - Custom plugin and scripts

Highlights include:

- Cloud Exchange is run in a Docker with a small compute footprint.
- CE can ingest, manage, and share millions of attack indicators.
- Out-of-the-box integration with CrowdStrike Falcon, VMware Carbon Black Cloud, SentinelOne, and ServiceNow.
- Define frequency of the updates for every connected pair.
- Open architecture to build and add integration plugins to handle sharing to and from CE between IT systems and/or scripts.
- Configurable Frequency of polling, timeouts, and sharing.
- CE Dashboard provides information on how often IOCs have been seen within the Enterprise, determining the scope of the attack surface.
- PEP for threat intelligence in real time, ensuring that enterprise users and systems do not access malicious sites, endangering organizational security posture.
-

PRINCIPLES

OVERVIEW

The Netskope Zero Trust architecture and principles follow the guidelines from NIST 800-207, DoD Zero Trust Reference Architecture, and DHS CISA Zero Trust maturity model. The essential basic tenets are:

Trust no one
Know your people and your devices

Validate identity at every step
Design systems assuming they are all compromised

Distrust everything
When a breach happens, you are as protected as you can be

Use Dynamic Access Controls
Access to services must be authenticated, authorized, and encrypted at all times and can be revoked during a session

Constantly evaluate risk

- Include context in risk decisions
- Monitor and log in every location possible
- Aggregate log, system, and user data

Right size protections

- Invest in defenses based on the classification of data.
- Spend more money defending the systems at greater risk.
    1. All data sources and computing services are considered resources.
    2. All communication is secured regardless of network location.
    3. Access to individual enterprise resources is granted on a per-session basis.
    4. Access to resources is determined by dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioral and environmental attributes.
    5. The enterprise monitors and measures the integrity and security posture of all owned and associated assets.
    6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed.
    7. The enterprise collects as much information as possible about the current state of assets, network infrastructure, and communications and uses it to improve its security posture.

Netskope delivers on these principles using a global SASE architecture that is FedRAMP authorized.

The zero trust architecture design and tenants exist with the following basic assumptions:

- The entire enterprise private network is not considered an implicit trust zone. Assets should always act as if an attacker is present on the enterprise network, and communication should be done in the most secure manner available (see tenet 2 above). This entails actions such as authenticating all connections and encrypting all traffic.
- Devices on the network may not be owned or configurable by the enterprise. Visitors and/or contracted services may include nonenterprise-owned assets that need network access to perform their role. This includes bring-your-own-device (BYOD) policies that allow enterprise subjects to use nonenterprise-owned devices to access enterprise resources.
- No resource is inherently trusted. Every asset must have its security posture evaluated via a PEP before a request is granted to an enterprise-owned resource (similar to tenet 6 above for assets as well as subjects). This evaluation should be continual for as long as the session lasts. Enterprise-owned devices may have artifacts that enable authentication and provide a confidence level higher than the same request coming from nonenterprise-owned devices. Subject credentials alone are insufficient for device authentication to an enterprise resource.
- Not all enterprise resources are on enterprise-owned infrastructure. Resources include remote enterprise subjects as well as cloud services. Enterprise-owned or -managed assets may need to utilize the local (i.e., nonenterprise) network for basic connectivity and network services (e.g., DNS resolution).
- Remote enterprise subjects and assets cannot fully trust their local network connection. Remote subjects should assume that the local (i.e., nonenterprise-owned) network is hostile. Assets should assume that all traffic is being monitored and potentially modified. All connection requests should be authenticated and authorized, and all communications should be done in the most secure manner possible (i.e., provide confidentiality, integrity protection, and source authentication). See the tenets of ZTA above.

- Assets and workflows moving between enterprise and nonenterprise infrastructure should have a consistent security policy and posture. Assets and workloads should retain their security posture when moving to or from enterprise-owned infrastructure. This includes devices that move from enterprise networks to nonenterprise networks (i.e., remote users). This also includes workloads migrating from on-premises data centers to nonenterprise cloud instances.

PILLARS AND CAPABILITIES

Netskope takes the Operational Concept and maps to six core pillars of Zero Trust as illustrated in the diagram below. These are User & Device, Network & Environment, Application & Workload, Data, Visibility & Analytics, and Automation & Orchestration.

Netskope Zero Trust - Pillars and Capabilities

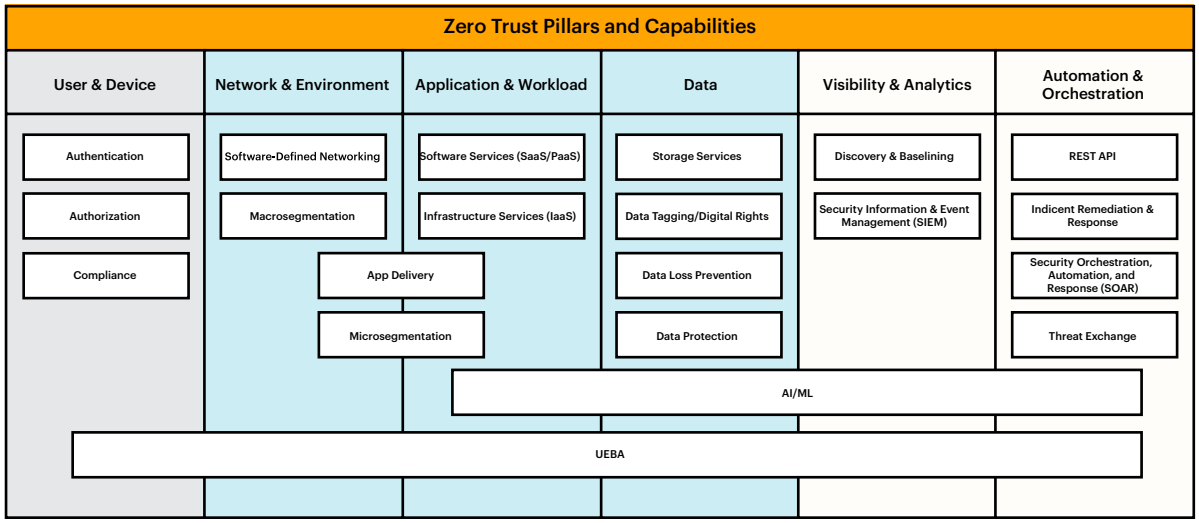

Each of the six pillars contain several key capabilities that map to one of the PEPs of the Netskope platform. These capabilities are:

User & Device

Authentication—Determines the user/entity who is looking to access the resource after a three-way bind of the user's identity allowing access to access resources.

Authorization—Determines the user/entity entitlement to access protected resources that could be internet or network bound.

Compliance—Determines user to application usage compliance, configuration of SaaS/IaaS/PaaS instances and enforces acceptable usage policies.

Network & Environment

Software-Defined Networking—Determines user network access to private applications and services inside the data center or virtual private cloud services.

Microsegmentation—Limits communication between devices, enforces policy before allowing connectivity at the Application level.

Macrosegmentation—Divides the network resources up into segments or trust zones allowing enforcement of policy to like groups. Macrosegmentation also works to provide isolation solutions including secure browsing solutions like RBI.

Application and Workload

SaaS/PaaS—Determines access to SaaS/PaaS offering, determines the instance, and enforces activity level control.

Macrosegmentation—Leverages Application Instance Awareness and Activity Understand to determine how a user should interact with an Application. Applications include SaaS Applications as well as Custom Applications defined within the Netskope platform.

Data

- Storage Services
- Data Tagging/Data Rights
- DLP
- Data Protection

Visibility and Analytics

- Discovery and Baselining
- SIEM

Automation and Orchestration

- RESTful API
- Incident Remediation and Response
- SOAR
- Threat Exchange

ZTA Implementation and Functional Capabilities

Identity, resource, and control capabilities delivering zero trust architecture are enforced from a single unified engine consisting of ML, AI, and traditional policy techniques for UEBA, Threat, DLP, Reporting, Logging, and Remediation. ML delivers capabilities supporting multiple use cases for identity, resource, and control capabilities from ransomware detection to user behavior analytics and real-time validation with risk-based enforcement of device identity and context for continuous validation of network and application access.

User behavior monitoring is performed by identifying and logging both normal behavior and abnormal deviations using risk thresholds; anomalies and risk scores are generated using 30+ detection methods including fact-based rules, heuristic detections and to establish user risk scores against threat vectors such as data destruction, breach detection, malware distribution, lateral movement, and impossible traveler use cases. Visibility for all user and device interactions is provided with transaction-level metadata with access rules that provide the ability to fully automate and orchestrate user profiling, group- based access, just in time /just enough access control using data context, device, and user security posture.

An automation engine with governance and remediation workflows to alerts, coaches, blocks access, or challenges a user with a step-up authentication request before processing a request by continuously analyzing seven dimensions of user activity (time of day, day of week, location, device, service, activity, and object) identifying deviations from the user's baseline behavior across one or more of these dimensions. Reporting and enforcement include aberrant usage patterns such as abnormal high traffic volumes, previously unseen user locations/devices/operating systems/activities within the applications, or any combination of these activities such as users using apps in the new location.

User Confidence Index (UCI) is integrated into a real-time policy engine that allows for policies based on IAM provided users and groups, MFA, and DLP profiles and other contextual details. RESTful API provides integrations to allow for UCI to automate movement of risky users into more restrictive groups. Both rules and behavior-based policies are used to identify threats with a rules-based engine that matches user activities against a set of predefined rules (policies). Rule-based anomalies would include events such as bulk downloads, bulk uploads, deletes, failed logins, proximity events, rare events, risky countries, and data exfiltration between agency and personal cloud instances matching against multiple activities across multiple cloud services.

Threat Protection within a zero trust framework is applied for both real-time inline or in near real time through API-based connections to cloud services. Threat capabilities include:

- Anti-malware engines, client traffic exploit protection (CTEP) – inbound browser IDS/IPS, and true file type analysis
- 40+ threat intel feeds, plus importing IOCs including malicious URLs and file hashes, STIX/TAXII
- Inline portable executable (PE) file ML static analysis
- Cloud Threat Exchange (CTE) enables threat intel sharing with EPP/EDR, SIEM, SOAR, etc.
- RBI for uncategorized and risky websites
- De-obfuscation and recursive file unpacking with support for 350+ families of installers, packers, and compressors
- Pre-execution analysis and heuristics for 3,500+ file format families, with 3,000+ static binary threat indicators for Windows, Mac OS, Linux, iOS, Android, firmware, Flash, PDF, and other document types
- Sandboxing for behavioral analysis of 30+ file types including executables, scripts, and MS 365 documents for Windows operating systems with the ability to defeat evasive techniques
- ML deep analysis to detect unknown threats, anomalies, and behaviors

Data context and data usage within instances of applications along with user source attributes are critical to a zero trust architecture. The DLP engine facilitates protections during user request to resources and limits privileges when risk thresholds exceed allowable limits. Usage context enforcement is based on underlying actions and access to an instance of an application while data context enforcement is performed by a cloud-based DLP engine supporting use cases like data exfiltration from sanctioned to unsanctioned services, real-time data control of sensitive content, closed-loop incident response workflows, orchestration for forensics, notifications, and approvals that surround DLP events. Netskope provides multiple regulatory compliance templates to identify sensitive data and more than 3,000 predefined data identifiers that can be used individually, or combined with other predefined or custom identifiers to detect specific data.

DLP capabilities include the following:

- Custom data identifiers, regular expressions, and dictionaries with natural language processing
- Rules engine supporting Boolean operators, proximity variables, density, occurrences, and record-based scans
- Fingerprinting with threshold analysis, optical character recognition, exact data matching
- Predefined and custom classifiers utilizing AI and ML to identify image, text, and screenshots likely to be containing sensitive information
- File properties, wildcard file names, file sizes, file hashes, file types, and embedded metadata such as labeling, rights management, password protection, and encryption
- ICAP/REST API integration for analysis and verdict; integration with DLP, EDR, DRM, SOAR, SIEM
- SMTP Proxy for email traffic from thick clients/certificate-pinned applications
- Application-based support for BYOD, managed, and unmanaged mobile devices

3.0 OPERATIONAL WORKFLOWS

INTRODUCTION

Operational workflows depict how ZTA is applied at a capability level. Each workflow maps to a capability in one of the pillars defined in Section 2.2.

3.2 USER/DEVICE

The User/Device workflow is a partial workflow that depicts the process of validating and authorizing the user or non-person entity (machine) prior to connecting to any resources. In our Zero Trust Architecture the state of the User/Device is available for use by any PEPs used by any other workflows. Note that almost all of the workflows will utilize and reference the User/Device workflow as well.

Netskope User and Device Zero Trust Implementation Flow Diagram

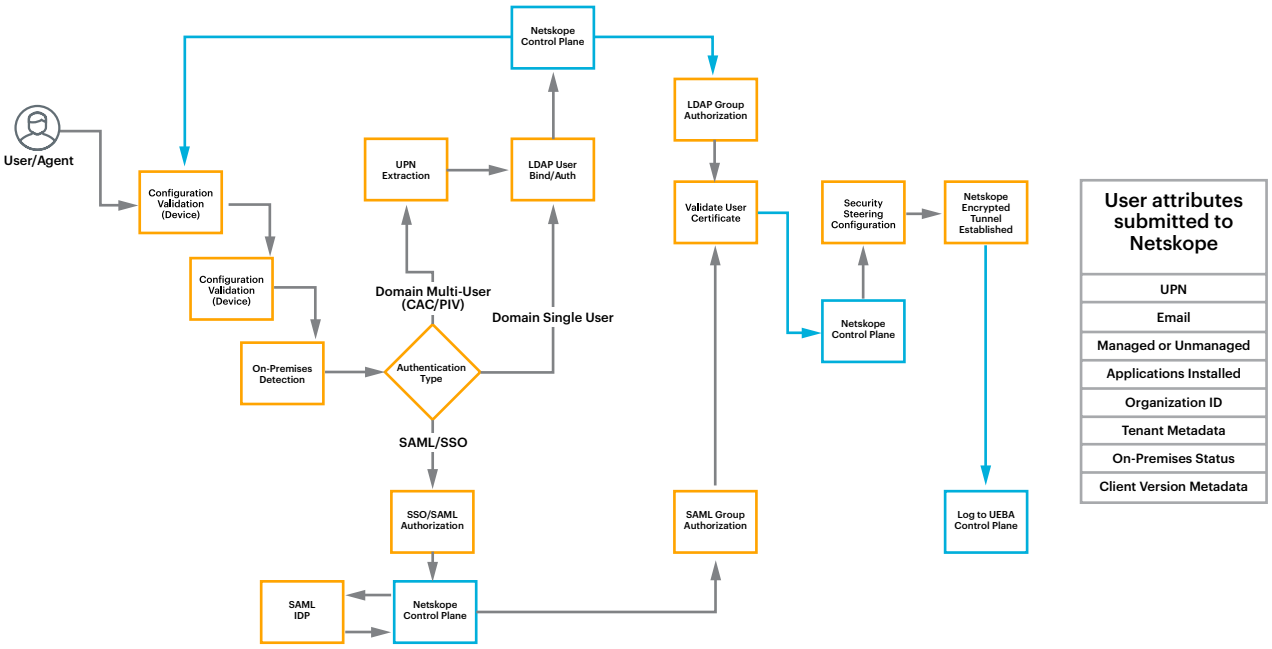

The User/Device utilizes an agent on the endpoint often referred to as a steering client. This agent generally provides three functions.

- Validation of the device
- Validation/authentication of the user
- Forward traffic to the Netskope Security Cloud

Validation of the device includes verification of the agent itself, ensuring that it has not been tampered with. It also involves checking the device certificate and determining if the device is a managed device or unmanaged (for future policy evaluation by the security cloud).

Validation and authentication of the user occurs next and is different based on the type of authentication. Generally the agent takes the user identification information and validates it against an authoritative authentication service and the Netskope Security Cloud. At this time the group membership is also determined.

The third function is the forwarding of traffic to the Netskope Security Cloud. The steering client will set up a transport layer security (TLS) tunnel using a client certificate generated during the authentication process. The client then forwards traffic to the Netskope Security Cloud via this secure TLS tunnel.

3.3 NETWORK/ENVIRONMENT

The Network/Environment Zero Trust Implementation pillar consists of multiple workflows for software-defined wide-area networking (SD-WAN) and Macrosegmentation. Application Delivery and Microsegmentation overlap with the Application and Workload pillar.

SD-WAN Workflow:

Netskope SD-WAN Zero Trust Implementation Workflow

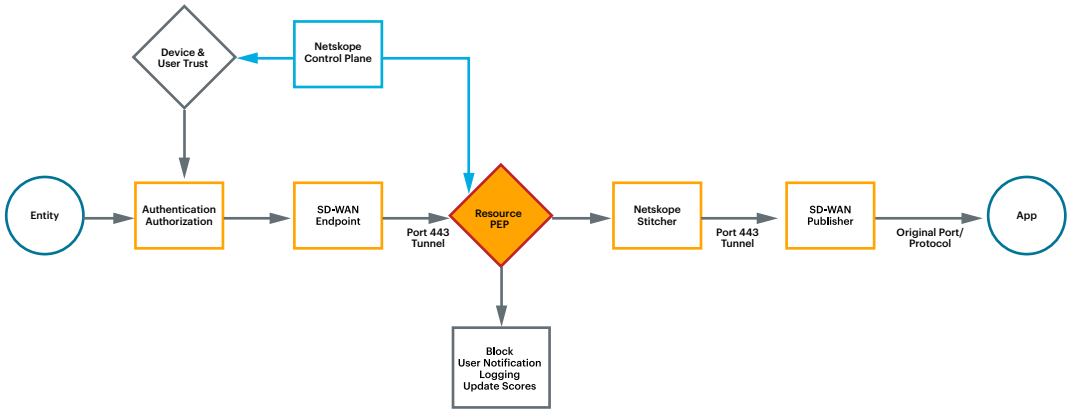

The SD-WAN workflow has three basic steps:

- It begins with the User/Device authentication workflow (see Section 3.2) and then continues as an SD-WAN endpoint. Each new session is forwarded over the TLS secure tunnel and is evaluated by the Netskope Security Cloud acting as a resource PEP.
- An SD-WAN publisher is pre-staged where the application resource is located. This publisher is authenticated to the Netskope Security Cloud via certification authentication and a secure tunnel is established to the Netskope service.
- When a user/entity requests access to the resource, the request is forwarded over the secure tunnel to the Netskope PEP and is evaluated against a configured policy. If the request is allowed, then it is forwarded via the Netskope stitcher to the SD-WAN publisher, which then forwards the request to the resource. The session is then allowed.

Macro Workflow:

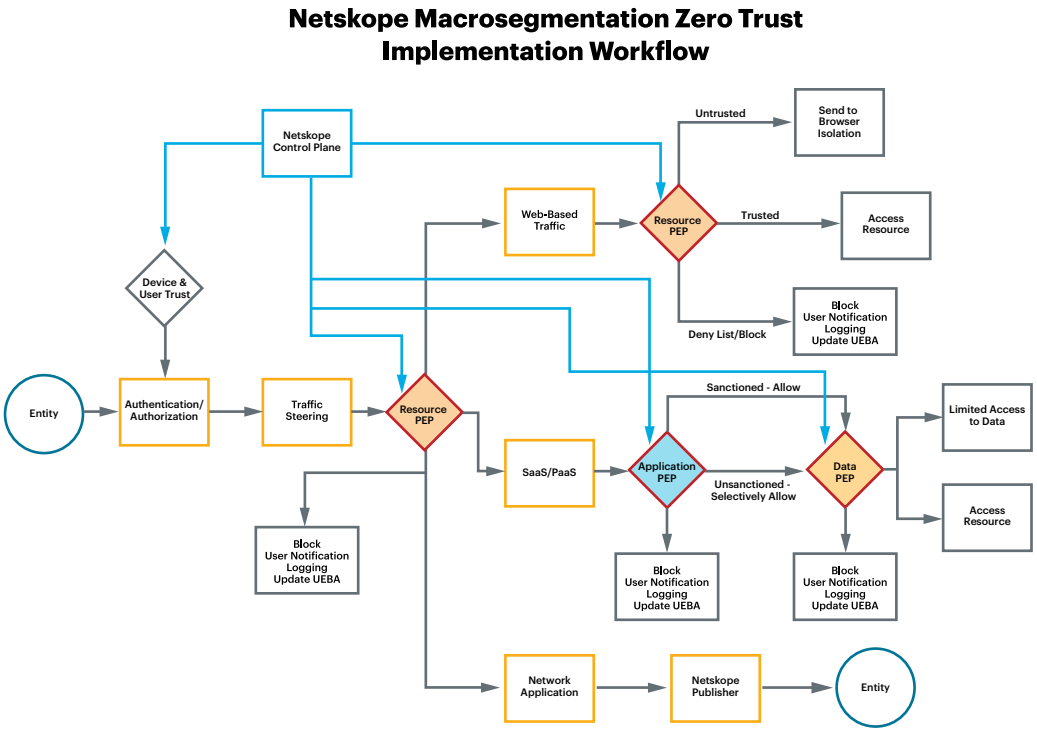

The Macrosegmentation workflow also begins with the User/Device authentication workflow (see Section 3.2) and then continues using the Netskope Security Cloud as a Resource PEP.

The Resource PEP evaluates traffic for policies that are SaaS/PaaS-based, network-based (L3/L4), or private application-based.

For network-based policies, the Resource PEP evaluates the L3/L4 rules and will apply the appropriate actions (Allow/Deny/Coach user). Policies are applied so that lateral movement by the client/endpoint is not possible.

For SaaS/PaaS data, the session is passed to the Application PEP for evaluation (and possibly the Data PEP) in Section 3.4. The Application/Data PEP workflows are then executed.

Each new session is forwarded over the TLS secure tunnel and is evaluated by the Netskope Security Cloud acting as a resource PEP.

3.4 APPLICATION & WORKLOAD

The Application workflow begins with the User/Device workflow (see Section 3.2) to establish the user identity and the device classification (managed, unmanaged, etc). The Netskope Security Cloud then decodes the session and determines if it is a cloud application or other type of session (web or non-web).

Once the cloud application is determined, the Application PEP is applied. The Application PEP can determine if access is allowed as well as what activities are being used (Create, Edit, View, Share, Upload/Download, Post, etc). Once the action is determined, the Application PEP will invoke the Data PEP (if necessary) and the Data Segmentation workflow will be executed.

3.5 DATA

The Data Segmentation Zero Trust Implementation workflow:

Netskope DLP/DRM Zero Trust Implementation Workflow

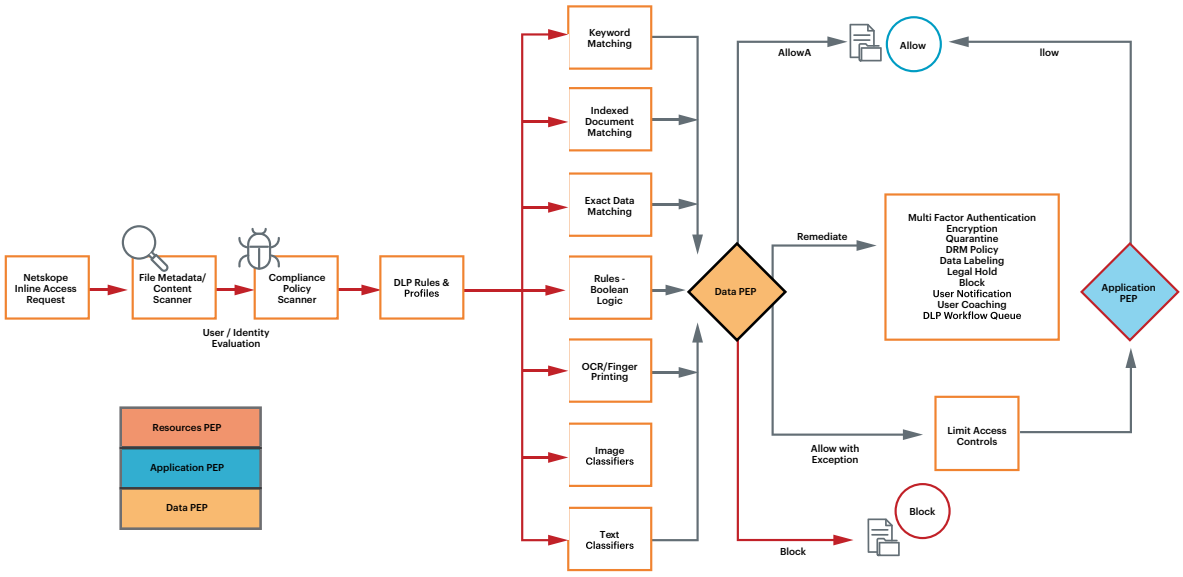

As with the other workflows, the Data Protection and DLP workflow begins with the User/Device authentication workflow (see Section 3.2). Each new session is forwarded over the TLS secure tunnel and is evaluated by the Netskope Security Cloud acting as a resource PEP and data PEP. The data is scanned for content and file types that can be recognized and evaluated by the DLP engines, which decode more than 1,500 file types and 3,000 data identifiers. Following the file identification, the compliance engine scans the data for compliance violations. Currently there are more than 40 predefined compliance templates. The data is then passed to the DLP engine, which is configured with preconfigured templates and/or custom templates for data protection violations.

The DLP engines will form a verdict and the data will either be blocked, remediated, allowed, or allowed and forwarded to the Application PEP. The Application PEP can then enforce application policy on this data (Share, View, Upload/Download, Edit, etc).

VISIBILITY, ANALYTICS, ORCHESTRATION, AND AUTOMATION

Visibility, Analytics, Orchestration, and Automation are contained within the control plane functions of Netskope. This is where we complete UEBA as well as tying in automation or third-party tools. User behavior monitoring is performed by identifying and logging both normal behavior and abnormal deviations using risk thresholds; anomalies and risk scores are generated using 30+ detection methods including fact-based rules, heuristic detections and to establish user risk scores against threat vectors such as data destruction, breach detection, malware distribution, lateral movement, and impossible traveler use cases. Visibility for all user and device interactions is provided with transaction-level metadata with access rules that provide the ability to fully automate and orchestrate user profiling,

group-based access, just in time/just enough access control using data context, device, and user security posture. An automation engine with governance and remediation workflows alerts, coaches, blocks access, or challenges a user with a step-up authentication request before processing a request by continuously analyzing seven dimensions of user activity (time of day, day of week, location, device, service, activity, and object) identifying deviations from the user's baseline behavior across one or more of these dimensions. Reporting and enforcement include aberrant usage patterns such as abnormal high traffic volumes, previously unseen user locations/devices/operating systems/activities within the applications, or any combination of these activities such as users using apps in the new location.

CONCLUSION

The idea of zero trust has created a lot of noise in recent years, but what we're seeing most recently with Executive Order (EO) 14028 is the much-needed planning, coordination, and preliminary action that's needed to make long-term critical improvements to federal security. As evidenced by both the recent OMB Memo on Protecting Critical Software Through Enhanced Security Measures and the Security Measures for "EO-Critical Software" Use Under EO 14028 released by NIST, what we're seeing is a genuine effort to reset the baseline of security in the federal government. It's important to remember the scale of this challenge holistically—redefining security across an organization with more than 4.2 million FTEs, not to mention contractors and support personnel, massively complex networks, and countless tools in use across each individual agency both for mission sustainment and security—there is no easy button or one-size-fits-all solution. So, while initial efforts and support documentation may not be comprehensive, and the adherence to the guidance may not eliminate the need to implement additional security measures, it does give agencies the freedom to pursue the most innovative technologies and modern security solutions on the market. We expect additional guidance from the White House, as well as appropriations and funding to support the changes required across the federal government. The agencies that will prove most successful in improving their security posture are those that start planning for zero trust adoption now and put together a long-term strategy for continuous modernization.