zscaler™ | Public Sector

AUTHORIZED StateRAMP

White Paper

The One True Zero Trust Platform for State and Local Government

Accelerate Secure Digital Transformation with the Zscaler Zero Trust Exchange

Introduction

Digital transformation isn't new, but for many organizations, the pandemic was a catalyst that accelerated adoption of digital processes and technology for serving constituents and equipping employees to work from anywhere. Moving applications from data centers to the cloud brings efficiencies and agility that allow state and local government to achieve secure IT modernization. Organizations can enable remote employees to work and collaborate more productively, protect against ransomware attacks, streamline operational processes, and revolutionize digital constituent experiences.

But *secure* digital transformation demands new ways of thinking about architecting and securing connectivity for people, apps, and data everywhere. For the past three decades, organizations have built complex, wide-area, hub-and-spoke networks to connect office-based users to applications residing in the data center. These networks relied on a stack of security appliances and firewalls to prevent anyone outside them from entering, but granted privileges to everyone within.

This model, known as castle-and-moat security, worked reasonably well when most applications lived in the data center and most users worked from the office. However, it wasn't scalable and became increasingly outdated as apps moved to the cloud and remote work adoption soared. Once users and applications moved off the network, security policies could only be enforced by backhauling traffic to a central data center for inspection. As a result, the more remote users there were, the more difficult it was to ensure a good user experience.

Today's cloud-first, geographically distributed computing ecosystems — in which users can connect from anywhere and access agency resources everywhere — need something different. In recent years, zero trust has been widely adopted for this purpose, coming to serve as a de-facto standard that enterprises and noncommercial entities (including government agencies) rely on to secure their users, apps, Internet of Things (IoT) devices, and workloads, regardless of which networks or cloud resources they're connecting to, where they're located, or how they're connecting.

Zero trust overcomes the inherent limitations of castle-and-moat security architectures, which inevitably make for an ever-expanding attack surface as IT environments grow more complex. In the old model, every internet-facing resource — from cloud applications and end user devices to virtual private networks (VPNs) and firewalls — could potentially be discovered and exploited. Once this initial compromise had been achieved, it was relatively simple for attackers to move laterally from one high-value target to the next within a corporate network where all traffic is trusted. And once the network was breached, attackers could take advantage of this trust to exfiltrate data from apps and data stores whether in the data center, SaaS apps, or public cloud.

The Evolution of Zero Trust

Although the idea of "de-perimeterizing" networks had been discussed for nearly two decades, the term "Zero Trust" was initially coined in a 2010 paper published by Forrester analyst John Kindervag. The paper argued that mere presence on a network should not be sufficient grounds to trust a user or device. As the key concept within this new way of thinking, zero trust quickly became a buzzword.

Shortly after the publication of Kindervag's paper, Gartner introduced the concept of Continuous Adaptive Risk and Trust Assessment (CARTA) that introduced the idea of granting access based on ongoing assessment of the environment, contextual information, and what users' responsibilities warrant. Over the years, CARTA evolved into Secure Access Service Edge (SASE), an architectural framework that brought together cloud-native security technologies (including ZTNA) and wide-access networking (WAN) capabilities to securely connect users, systems, and endpoints to applications and services anywhere. By 2021, Gartner divided the SASE

market category into two subcategories: WAN and the Security Service Edge (SSE), in which SSE represented a converged set of security services delivered from a unified cloud platform. SSE enabled an organization to enforce zero trust policies everywhere, even if users are off network, to meet the dynamic secure access needs of modern digital enterprises.

Zero trust became a national standard within the US with the 2000 release of the National Institute of Standards and Technologies (NIST)'s Special Publication 800-207. This document explicitly defined zero trust as an architectural model with the underlying principle that "no implicit trust [should be] granted to assets or user accounts based solely on their physical or network location (i.e., local area networks vs. the internet) or based on asset ownership (enterprise or personally owned)." In other words "never trust, always verify." Finally, in 2022, a new regulatory mandate by the US government to build security architectures based on NIST principles led to zero trust becoming the default security paradigm for protecting applications, traffic, users, workloads, and devices in modern computing environments.

However, confusion in the marketplace still persists about what is — and isn't — a true zero trust solution. This confusion has been nurtured by perimeter-based security vendors who stand to lose market share when customers realize that their solutions do not include the core components necessary to build a zero trust architecture, and cannot perform the requisite functions to secure it. It's important to note that a zero trust architecture is fundamentally opposite of a network security-based one. It's impossible to implement zero trust on a routable network by using firewalls and VPNs. Zero trust requires verifying identity and context before allowing access to resources, not connecting through a network.

To cut through the noise — and help security and IT teams understand how to build security architectures that meet the needs of modern organizations — we've put together this guide. It explains what zero trust architectural principles actually are, and how a solution like the Zscaler Zero Trust Exchange™ makes it easy to implement them.

What is a Zero Trust Architecture?

Zero trust begins with the assumption that everything on the network is hostile or compromised, and access to an application is only granted after user identity, device posture, and agency context have been verified and policy checks enforced. In this model, all traffic must be logged and inspected — requiring a degree of visibility that cannot be achieved with traditional security controls.

A zero trust architecture is expressly designed to minimize the attack surface, prevent lateral movement of threats, and lower breach risks. It's best implemented with a proxy-based architecture that connects users directly to applications instead of the network, and that makes it possible for additional controls to be applied before connections are permitted or blocked.

A true zero trust architecture differs from traditional architecture in three ways:

1. **Terminates every connection.** This differs from the pass-through approach employed within technologies like firewalls that inspect files as they are delivered. If a malicious file is detected, it's often too late by the time the alert is delivered. Terminating every connection instead allows for inline inspection of all traffic. This prevents ransomware, malware, and malicious traffic from ever reaching their destination.

2. **Protects data using granular, context-based policies.** Zero trust architecture should verify identity and context — user identity, device identity, device location, type of content, and the application to which access is being requested — before granting access. These policies should be adaptive so that access privileges can continually be reassessed as conditions or user behaviors change.

3. **Eliminates the attack surface.** Once connected, anyone who is on a traditional network can see all other nodes of that network. By contrast, in a zero trust architecture, users have visibility of — and the ability to connect to — only the resources they are authorized to access. Nothing more. Direct user-to-app and app-to-app connections eliminate the risk of lateral movement and prevent compromised devices from serving as a source of infection that spreads to other resources.

One True Zero: How the Zero Trust Exchange Provides the Key Building Blocks for a Zero Trust Architecture

Zscaler has been a leader in zero trust security for over a decade. To help organizations advance their digital transformation securely, Zscaler created the Zero Trust Exchange, an integrated cloud-native cybersecurity platform founded on the principle of least-privileged access, and the idea that no user, workload, or device is inherently trustworthy. The platform grants access based on identity and contextual information such as device type, location, application, and content, brokering a secure direct connection between an application and a user, workload, or device — over any network, from anywhere.

The Zero Trust Exchange is an integrated platform of services that acts as an intelligent switchboard. Its unique proxy architecture can guarantee the enforcement of zero trust policies, regardless of location. This approach treats all communications as potentially hostile, blocking all of them until they can be validated according to identity-based policies. Every communication that flows through the Zero Trust Exchange is subject to a series of controls before a connection is established.

There are seven essential elements in a zero trust architecture.[1] In turn, these elements can be grouped into three categories:

1. **Verify Identity and Context**
   Whenever a user, device, or workload requests a connection — irrespective of the underlying network — the Zero Trust Exchange first terminates the connection and verifies identity and context by understanding the "who, what, and where" of the request.

2. **Control Risk**
   Once the identity and context of the requesting entity have been verified and segmentation rules applied, the Zero Trust Exchange then evaluates the risk associated with the connection request. It also inspects the traffic for cyberthreats and sensitive data.

3. **Enforce Policy**
   Finally, the platform uses the outputs of the previous elements to enforce policy on a per-session basis, and ultimately determines whether to conditionally allow or conditionally block the connection. If the entity is allowed to connect, the platform ensures that its connection to the internet resource, SaaS app, or IaaS/PaaS environment is a secure one.

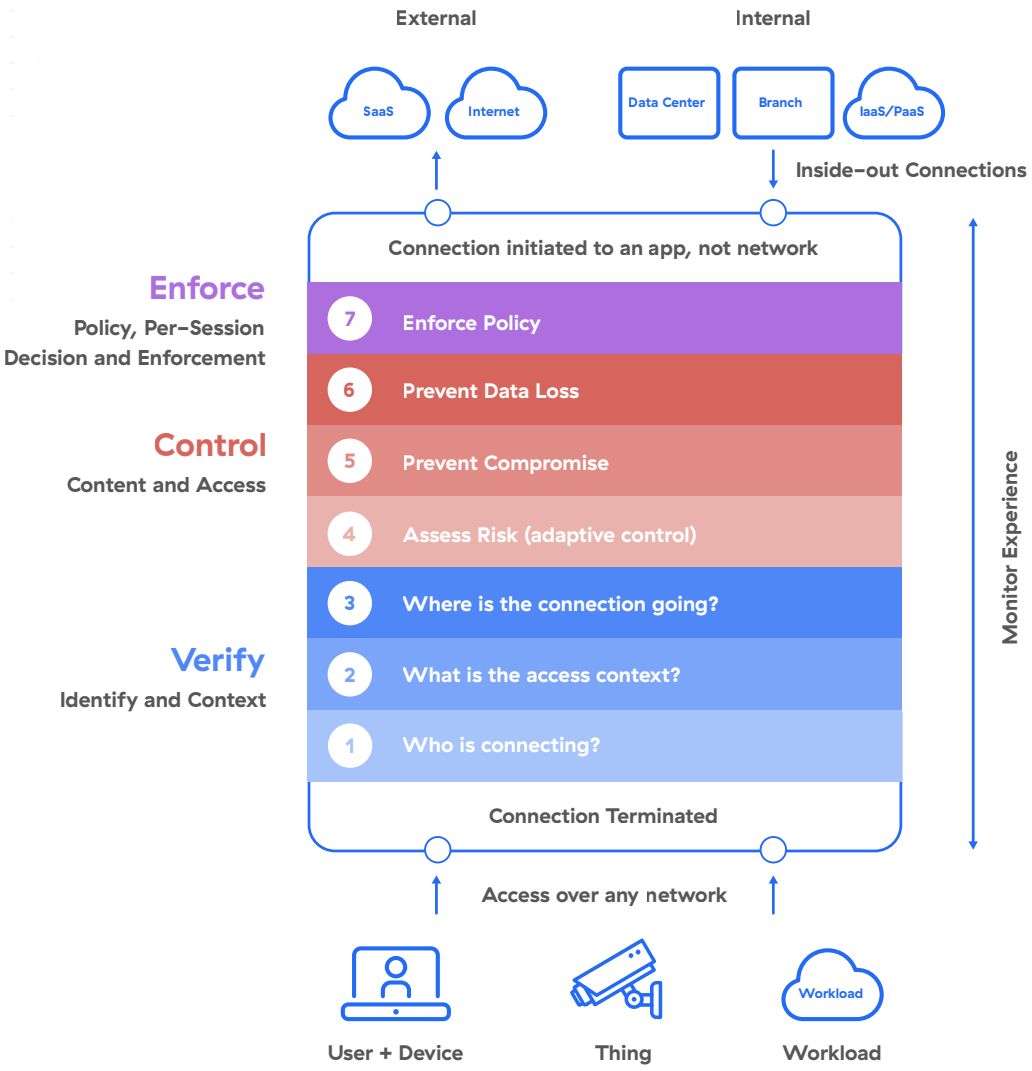

Figure 1: The seven elements of zero trust architecture.

Each of the seven elements within the three core categories feeds into the others, creating a dynamic decision tree that's used to verify every connection, every time. In the Zero Trust Exchange, these policy decisions and enforcements are achieved without compromising user experience: the platform monitors performance and diagnoses issues to ensure that security doesn't place an undue burden on end users.

Let's take a closer look into each of the seven elements.

Verify

At this stage in the connection process, the Zero Trust Exchange ascertains:

- Who is connecting
- What the access context is
- Where the connection is going

1. **Who is connecting:** The Zero Trust Exchange typically establishes who is connecting by leveraging an extensive set of integrations with identity providers (IdPs). The IdP platform will validate the connecting entity's

identity through authentication (ideally MFA or stronger), and the Zero Trust Exchange will then consume identity values, certificates, and shared details that the IdP provides. To identify devices that cannot be authenticated by an IdP, such as IoT or operational technology (OT) devices, the Zero Trust Exchange can leverage network location information. It can also identify individual workloads on the basis of characteristics, process functions, or other factors.

2. **What the access context is:** When the Zero Trust Exchange considers access context, it evaluates granular detail around the context of the connection. Access context validation is based on identity and profile information. Instead of merely showing that someone is an employee, this information will reveal more about that person's intent and permissions, allowing for a more nuanced access decision.

   Additional access context attributes can include:

   - Device location
   - Request time
   - User's job, role and responsibilities
   - How the access request is being made (e.g., does it fall within normal patterns?)
   - Device identity (personal vs. enterprise, managed vs. unmanaged)

3. **Where the connection is going:** After the Zero Trust Exchange has investigated the identity that's requesting the connection and learned about the context of the request, it will then ask where the connection is going. The Zero Trust Exchange will identify the app that's being requested, looking at its function, location, known risks and issues, and relation to the identity of the access requester.

   Is the app known or unknown? Is it available on the public internet?

   Because the Zero Trust Exchange works within a proxy-based architecture, it can focus on the context of the destination application rather than just its IP address. This makes it possible to achieve granular user-to-application segmentation that's based on identity-aware policies. And every application can be evaluated individually, so that enterprise resource planning (ERP) software will be treated differently than YouTube.

   This way, the Zero Trust Exchange can deliver true least-privileged access, enforcing specific policies for individual applications — including mission-critical workloads. It can also learn as it goes along, deriving insights from typical user-to-workload traffic flows, and iterating upon the policies in response. And it's intelligent enough to be able to assess the risk levels associated with unknown apps that are available for consumption on the open internet.

Control

In general, the decisions and enforcement actions that fall within this category are concerned with understanding risk and inspecting content. This degree of control cannot readily be achieved within conventional firewall-based architectures where it's not possible to inspect full content without adding on additional layers of services.

At this stage in the connection process, the Zero Trust Exchange will:

- Assess risk by applying adaptive controls
- Prevent compromise
- Prevent data loss

4. **Assess risk (with adaptive controls):** The Zero Trust Exchange dynamically assesses risk on an ongoing basis. It leverages signals that are continuously updated to ensure that the risk scores remain accurate. These risk scores then feed a decision engine that determines whether access should be granted — or should continue to be granted. This makes it possible to enforce risk-based access decisions over the lifetime of a connection. A change in user/device posture or behavior will trigger an update to the access decision in real time.

   Risk scores take into account factors like behavior (an SQL client should communicate with an SQL server, not an unrecognized server in an unfamiliar geographical location) as well as insights from third-party tools like EDR, security information and event management (SIEM) and/or security orchestration, automation and response (SOAR) solutions.

   Risk scoring can also consider things like industry-related risks, risk distribution across an organization (is IT riskier than sales?), and location-related risks. The Zero Trust Exchange's proprietary algorithm calculates a risk score on the basis of all of these factors. Its goal is to enable greater control and visibility at the point of policy enforcement. It can also correlate insights from cloud-native security engines to bring to light hidden risks posed by misconfigurations, threats, and vulnerabilities within the cloud stack. In addition, it can adapt scoring based on the latest findings from ThreatLabZ, Zscaler's embedded research team of security experts and network engineers, who analyze threats across the entire Zscaler security cloud as well as the global threat landscape.

5. **Prevent compromise:** The Zero Trust Exchange provides full inline content inspection to prevent compromise. The majority of internet traffic (including attack traffic) is now SLS/TLS-encrypted, so SLS/TLS inspection has become a must. It's become the only way to protect against initial attacks or stop the exfiltration of encrypted data. This inspection is performed in a way that's based on agency risk and application type in order to protect end users' right to privacy.

   The Zero Trust Exchange architecture is purpose-built for performance and scale. Its edge-delivered zero trust solution scans all content in a single pass without copying packets and adding latency. Plus, the cloud-native forward proxy architecture allows for the use of technologies like sandboxing and browser isolation, where threats can be quarantined and pixels are streamed to the end user instead of an actual web page. And it can apply the right control for the right application in ways that are intelligent (for instance, there's no need to inspect video stream content from Zoom or Teams).

6. **Prevent data loss:** The Zero Trust Exchange also delivers data loss prevention capabilities. By monitoring what's being sent out to the open internet — SaaS or internal applications hosted in public clouds — it can stop the exfiltration of sensitive data or intellectual property. These capabilities can also protect against accidental data loss and oversharing, as well as misconfigurations (e.g., of cloud storage).

When it comes to data loss prevention, the Zero Trust Exchange is highly capable. Its AI-driven inspection engine can enforce policies based on predefined dictionaries or engines for compliance purposes. Or users can build custom dictionaries based on enterprise-specific keywords and patterns. It's also able to perform advanced data classification. The Zero Trust Exchange enforces data protection policies at the edge, meaning that it happens close to the user, instead of in a way that's centralized, which would introduce latency. Enforcement is constant, and policies are uniformly applied. The Zero Trust Exchange is also able to scan APIs from SaaS providers to protect data at rest in SaaS apps, and it can scan data in motion to protect it there as well. It offers a granular set of out-of-band controls based on cloud app definitions, file type controls, and risk attributes. This means it can define acceptable file types, determine which cloud apps are allowed (and which are not), and even block sensitive data found in images, screenshots, or documents.

Enforce

During the previous two stages, the Zero Trust Exchange will have made a risk assessment based on identity and context. This is not a static operation that's done only once, but an ongoing, dynamic process.

7. **Enforce policy:** In the enforce stage, the policy engine considers the output of the risk assessment on a per-session basis, to determine whether to issue a "conditional allow" or "conditional block" decision, though deny decisions can be made prior to this stage.

   "Conditional allow" grants access to the application, but the platform can still deliver additional controls such as browser isolation, content inspection, and warnings. Within this category, there are several different types of controls that can be enforced. These include:

   - **Warn and allow:** access is granted, but the user is presented with a warning that the destination's risk is unclear.
   - **Prioritize:** with this bandwidth control function, dedicated network links are established to preserve access to business-critical applications, deprioritizing (for instance) social or streaming media.
   - **Isolate:** this control renders the requested content as a stream of pixels instead of a full web page, removing the risk of data leakage or active threats.
   - **Steer:** this sends the traffic to a non-standard destination.
   - **Quarantine and allow:** this function uses cloud sandboxing to "detonate" potentially harmful content, granting access only if the material turns out to be benign.

   "Conditional block" takes place if an access request does not meet the conditions for which it's evaluated in the previous steps. Included blocking controls include simple blocking, deploying decoy assets as a detection strategy for active threats, and blocking plus quarantining. The Zero Trust Exchange's ability to deliver multilayered policy enforcement means that it can provide powerful but granular controls and make precise but nuanced decisions. Enterprises can build various levels of policy enforcement based on the outcomes of the previous six elements, and multiple policies can be applied within the span of a single session. This makes it possible to create desired agency

outcomes, mitigate risks, and enforce a robust security posture.

The Zero Trust Exchange: The One True Zero Trust Platform

The Zscaler Zero Trust Exchange provides a truly comprehensive set of cyberthreat protection and connectivity capabilities — all within a single solution. It eliminates the attack surface and prevents data loss, using an architecture that was purpose-built to enforce zero trust policies for all users, workloads, and devices in an enterprise environment.

The Zero Trust Exchange:

- Simplifies and automates the process of connecting workloads from any location to any destination (whether that's in the public cloud or a private data center). Unlike VPNs, which increase risk by connecting workloads to networks, the Zero Trust Exchange extends connectivity only to workloads that require it —according to agency policies. This approach makes it nearly impossible for malware to spread, or for attackers to move laterally across a network.
- Inspects all content, including SSL-encrypted content, to apply powerful cyberthreat defense and data protection controls.
- Makes applications invisible from the internet through its unique proxy-based architecture, which eliminates the external attack surface. It verifies an entity's identity and determines the context of the access request before granting (or conditionally blocking) access.
- Integrates the ecosystem of security and IT application providers required to support global deployments, including identity, endpoint detection & response (EDR), security and IT operations tools, SaaS applications, SD-WAN, and more.

This architecture is purpose-built to deliver superior user experiences at scale. Unlike legacy architectures in which all traffic is backhauled to a data center for processing, the Zero Trust Exchange delivers direct connectivity to any cloud or internet destination. Traffic is intelligently routed to the nearest Zscaler location (among 150 globally-distributed data centers that have peering relationships with major cloud providers like AWS and Microsoft Azure) ensuring the shortest path of communication, no matter where applications are hosted. All content is scanned within a single pass, so there's no need to copy packets (adding latency).

The exchange also eliminates the operational inefficiencies that come with complexity. All SaaS, internet and private applications can be secured within a single platform, removing the need to maintain multiple hardware-based or virtual security appliances. A unified, cloud-native zero trust platform is quick to configure, easy to manage, and much more scalable than perimeter-based solutions. Inline policy enforcement also greatly simplifies the process of translating agency rules into network policies.

This approach reduces the costs associated with digital transformation. Security teams no longer need to budget and plan for the procurement of firewalls, VPN solutions, or costly MPLS networks with complex routing, switching, and network segmentation needs. The Zero Trust Exchange also reduces deployment timelines from months to days, while accelerating the organization's ability to detect and prevent data breaches that could cost millions.

The Zero Trust Exchange is a StateRAMP Authorized comprehensive platform offering a breadth of functionality that eliminates the need for point products:

- **Cyberthreat protection:** Users, workloads, and devices sit behind the exchange, which renders them invisible from the public internet. Because they can't be discovered, there's no attack surface to exploit. Plus, the exchange inspects all traffic, using ML-based advanced threat protection to prevent compromise.
- **Data protection:** The Zero Trust Exchange secures data across IaaS/PaaS, SaaS, email and endpoints, preventing data loss through advanced data classification and inline inspection of outbound traffic.
- **Zero trust connectivity:** Traditional networks are routable, which means that once you gain access, you can get anywhere on them. The Zero Trust Exchange fundamentally alters the nature of network connectivity by connecting users to apps rather than networks. This makes lateral movement impossible.
- **Digital experience management:** The Zero Trust Exchange monitors end-to-end user experience from the endpoint to the application. Its inline AI engine can pinpoint the root causes of issues, enabling IT to resolve them proactively.

Conclusion

The Zero Trust Exchange makes it possible to achieve a zero trust architecture that's seamless, secure, *and* cost-effective — without needing to make compromises. This solution is unique in the market because its purpose-built architecture enables organizations to:

- Eliminate the attack surface. Because applications are invisible from the public internet, there's zero external attack surface.
- Deliver superior end user experiences at scale. With 150 data centers distributed globally and peering relationships with major cloud providers, the Zero Trust Exchange intelligently manages and optimizes direct connections to any cloud or internet destination, with no need to backhaul traffic. This ensures low latency and top-notch performance.
- Eliminate operational efficiencies due to complexity, because of its all-in-one, single platform approach.
- Reduce the cost of digital transformation.
- Eliminate multiple security appliances and point products.

The Zero Trust Exchange aligns with the architectural requirements of well-known zero trust standards, including NIST 800-207 and Gartner's SASE and SSE frameworks. Proven at global scale, the Zero Trust Exchange is the world's largest security cloud, processing over 250 billion transactions per day. This scale means its ML engines are trained on more data than other platforms, improving the accuracy of threat detection, making it possible to stop even today's most sophisticated threats, while ensuring that end users can collaborate seamlessly and stay productive. As a comprehensive cloud platform, the Zero Trust Exchange enables organizations of all sizes to achieve a fast, reliable, and easy-to-manage zero trust architecture, while reducing the cost and complexity of point products.

Source: 1. Nathan Howe, Sanjit Ganguli, and Gerard Festa, Seven Elements of Highly Successful Zero Trust Architectures, Second Edition, August 2022.

+1 408.533.0288 Zscaler, Inc. (HQ) • 120 Holger Way • San Jose, CA 95134 zscaler.com