



Zero Trust, Maximum Resilience

A practical guide to today's “never trust, always verify”
security approach





Table of Contents

Why zero trust security's time is now	3
Who can benefit from zero trust	4
Selling the advantages	5
Key concept: Trusted device SSO	5
Key concept: Automated identity governance	6
Key concept: Phishing-resistant MFA and passwordless	6
A modeled approach	7
5 pillars of success	8
Cultural realities of zero trust	14
6 indicators of zero trust success	15
6 best practices in architecting for zero trust	16
Cloud providers and zero trust	20
How Oracle helps	21



Why zero trust security's time is now

By Lorna Garey
Senior Writer

Looking for a more proactive approach to security? The Pentagon is fast-tracking adoption of its zero trust architecture, with an initial 2027 deadline being moved forward a full 12 months. One industry watcher surveyed about 650 IT and cybersecurity professionals in a mix of industries and found that slightly more than one-third have already implemented a zero trust strategy, with an additional 47% planning to within a year.

If you're not among them, it may be time to explore what's driving this major strategic shift.

Zero trust security, also known as perimeterless security, uses frequent authentication and authorization, encryption, and granular segmentation to protect assets. No entity—person, device, or application, inside or outside the network—is trusted by default. Because organizations assume their systems are already compromised, they constantly monitor and reauthenticate users, services, and devices. They grant access based on data—the context of the request, the level of trust, and the sensitivity of the asset—and with only the permissions necessary to accomplish the task. Segmentation limits what an entity can access without reauthenticating, keeping an attacker from freely traversing the network and shortening the time to respond to a breach.

If that sounds like a dramatic change from the traditional “castle and moat” model, it is. And while a phased approach is common, moving to zero trust takes a lot of work, resources, and outreach. But it's arguably necessary to stay abreast of the threats most organizations will face in 2025 and beyond. We liken security controls to the brakes on a car—they're there to let you go fast with the assurance you can tap them as needed. Without the security they believe can protect the enterprise, IT leaders may default to a go-slow stance when it comes to green-lighting new technologies and using sensitive data.

Who can benefit from zero trust

A zero trust shop positions security controls strategically across the network, not only at the perimeter. That makes a zero trust architecture especially effective for organizations that run workloads in the cloud, share data with trusted partners, and support remote workers and locations and a diversity of devices. Still, zero trust isn't only for cloud workloads. Success comes with standardizing these principles across on-premises, hybrid, and cloud environments.

The Pentagon is an example of an extremely complex ecosystem that's all in on zero trust. David McKeown, the US Defense Department's deputy chief information officer for cybersecurity, [is on record](#) saying his team is taking a hybrid implementation approach that includes integrating new tools and capabilities; adopting commercial cloud solutions that have zero trust capabilities built in; and using purpose-built, on-premises private clouds.

As the DoD demonstrates, zero trust security isn't a product or service that can be switched on with the push of a button, though the right technology is critical. Rather, it's a way of working that may demand cultural, philosophical, and workflow adjustments.

How Zero Trust Works

Assume no entity is trustworthy, and grant only necessary access with the least privilege possible. Segment the network to restrain any bad actor who may get in.



People



Services



Devices

Authentication:

Who are you? Use strong passwords and multifactor identification.



Context:

Where are you, and how secure is your device? Limit access based on contextual data.



Least privilege:

What's the minimum access needed to fulfill your request? Make time- and resource-limited access grants.

Microsegmentation of the network to the extent possible is a cornerstone of zero trust.



Selling the advantages

Business leaders may ask, “Why now?” for zero trust. A confluence of reasons. Companies have more options for authentication, thanks to software, services, and hardware that increasingly support biometrics and sign-on methods beyond passwords.

And generative AI can help enable a zero trust architecture in a few ways:

- GenAI-powered automated threat detection systems can analyze huge volumes of network traffic and system logs to help identify anomalies and potential threats—and where authorized, stop them in real time.
- AI can establish baselines for normal user behavior and analyze patterns to help detect anomalous activity that might indicate, for example, someone trying to log on with stolen credentials or download valuable data.
- AI agents can automate provisioning and deprovisioning, reducing the risk of human errors.

Moreover, AI agents can help make sure each system or user is granted only the least access privilege necessary to get the job done, another tenet of zero trust. These factors combine to make 2025 the year the zero trust model gets more realistic for more organizations.

Key concept: Trusted device SSO

Trusted device single sign-on (SSO) allows users to access multiple applications and services with a single authentication, provided they’re using a trusted device. This can simplify the login process and enhance security by reducing the number of passwords or other forms of authentication people need to use.

To get started, register trusted devices. This often involves verifying a device’s specs and security. Once a device is trusted, its user may be able to log in with fewer additional authentication factors. However, IT should continuously monitor trusted devices for changes in their security postures. Trusted status may be revoked or, if a device is deemed risky, maybe due to its location or outdated software, additional authentication or actions to come into compliance can be required.

Key concept: Automated identity governance

Identity governance is the process of managing digital identities and access rights across an organization. It involves automating identity lifecycle management, enforcing access policies and separation of duties with strong access controls, and monitoring user activities in line with the zero trust approach of granting only the minimum necessary privilege.

The first bucket, identity lifecycle management, requires the ability to automatically create, modify, and delete accounts; assign roles to users based on their job functions and responsibilities based on separation of duties policies; and define granular access controls based on relevant attributes.

Enforcing access policies for user accounts and sensitive data revolves around granting only the minimum privileges to perform tasks, then monitoring and regularly reviewing and auditing permissions and activities to help detect anomalies and potential threats and check on compliance with industry regulations and internal policies.

Key concept: Phishing-resistant MFA and passwordless

Phishing-resistant multifactor authentication (MFA) methods look to keep attackers from intercepting one-time password (OTP) codes, while passwordless access is a security approach that supplants or supplements passwords with stronger authentication methods. People may authenticate using biometric factors, such as a fingerprint or facial recognition, in addition to a trusted device, such as a smartphone. Companies may also issue physical security keys that provide strong authentication without relying on passwords. And IT can assess the risk of each authentication attempt—such as from a user who's traveling or hasn't regularly used the system—and require additional verification steps if deemed necessary.

Besides lowering the risk of password-related breaches, companies could save money from fewer password resets and account lockouts. AI can help here, too, by regularly affirming a user's identity through passive facial scans or possibly voice recognition, further improving security.





A modeled approach

Companies don't need to reinvent the wheel. Multiple agencies have published freely available zero trust model architectures and technical guidance.

Cybersecurity and Infrastructure Security Agency Zero Trust Maturity Model

CISA is a US agency responsible for protecting critical infrastructure and promoting cybersecurity. Version 2.0 of its Zero Trust Maturity Model, or ZTMM, comprises five pillars plus several overarching principles and aims to be adaptable for a rapidly evolving business, security, and technology landscape.

🔗 [CISA](#)

Defense Information Systems Agency Zero Trust Reference Architecture

This document embeds zero trust security principles throughout the architecture and is used by the DoD and some other government agencies.

🔗 [DISA](#)

National Cyber Security Centre Zero Trust Architecture Design Principles

This flexible framework outlines eight principles to help implement a zero trust network architecture. The NCSC is a UK government entity that provides advice and guidance to businesses, other agencies, and individuals.

🔗 [NCSC](#)

National Institute of Standards and Technology Zero Trust Architecture

This framework, Special Publication 800-207, provides a comprehensive approach to implementing zero trust covering identity, devices, networks, and applications. NIST, a nonregulatory agency within the US Department of Commerce, promotes innovation and industrial competitiveness.


🔗 [NIST](#)

5 pillars of success

Many of these zero trust models are based on pillars and principles. The US CISA ZTMM, for example, is based on two core premises:

- The goal is to keep unauthorized parties away from data and services by making access control enforcement as granular as possible.
- No user or asset is to be implicitly trusted because it's assumed that a breach has already occurred or will occur, and, therefore, no entity should be granted access to sensitive information based on a single verification done at the perimeter. Instead, each user, device, and application or service must be known and continually verified with each transaction and the user's device health and behavior continually assessed and monitored.

The ZTMM also charts maturity along each pillar, from a traditional to an optimal approach, so teams can assess how they're doing in each area. Let's look briefly at the ZTMM pillars and what each prescribes.





Pillar: Identity

Zero trust security is predicated on strong identity and access management (IAM), underpinned by an identity governance and administration service that provides visibility across both cloud and on-premises systems. Identity can be associated with a human, an application, or a device and is managed with an IAM platform that confirms a user is who she says she is and is using a known device. Although a capable IAM platform can serve as a main identity repository, many organizations have multiple identity management systems, all of which need to be managed to achieve a zero trust architecture.

AI can help by generating analytics-driven access reviews that provide insights into access permissions and cloud infrastructure policies and call out potential anomalies.

Other enablers for this pillar include role-based access control (RBAC) supplemented with attribute-based access control (ABAC) and policy-based access control (PBAC) technologies. An RBAC system assigns permissions to users based on their roles within an organization and within an application. It may also dictate what data a user in a given role can access and how they can act on it.

ABAC is a more granular approach that considers additional resource and user attributes, such as time of day, location, and data sensitivity. PBAC is an even more advanced and flexible system that allows organizations to define complex policies based on rules that specify who can access what resources under what conditions. Policy enforcement points, or PEPs, help regulate access by evaluating requests against those rules. For example, by policy, an HR pro might be able to see only certain data within a CRM, such as sales performance stats.

Although the Extensible Access Control Markup Language, or XACML, standard defines security policies, companies should check that the access control platform they select can work with their critical applications. Open Policy Agent, an open source, flexible policy engine that unifies policy enforcement across different layers of an organization's infrastructure, is becoming a popular option as an add-on to XACML.

Strong authentication is another important element of this pillar. Someone may begin with a username and password—something they know—which is the first factor, and then be prompted to provide a second factor they have, possibly a code from a registered device or a fingerprint.

Just-in-time (JIT) access provisioning grants the minimum level of access required to perform a specific task and only for the duration needed. Access is provisioned on demand and automatically revoked when it's no longer needed.





Pillar: Devices

Devices aren't just smartphones and PCs but also systems used at the core of an owned or cloud data center. Effective device management for zero trust requires a dynamic inventory of assets, including hardware, software, firmware, and configurations, per our example CISA model. In general, it's good practice to grant permissions conservatively and limit how long a device can remain connected to the server based on inactivity or a set session duration.

Enabling technologies for this pillar start with endpoint security and mobile device management (MDM) systems because the “never trust, always verify” principle means that every access request must be thoroughly authenticated and explicitly authorized. Endpoint security solutions monitor devices, including desktops, laptops, and servers, to check that they comply with security policies, while MDM helps IT manage mobile devices. Together, they help create a robust security framework that predicates access on identity and contextual information.

Endpoint detection and response (EDR) systems help detect and respond to threats on devices, including laptops, desktops, and servers. EDR depends on advanced analytics and machine learning to continuously collect and analyze data that can, for example, identify potential misuse of a device or a malware infection.



Pillar: Networks

Zero trust mandates a shift away from traditional perimeter-focused security. In the CISA example, IT teams need to granularly manage internal traffic flows, isolate hosts, enforce encryption, segment activity, position security controls closer to assets, and achieve enterprisewide network visibility. It's a big challenge that depends on microsegmentation—dividing the network into contained zones, controlling movement between them, and applying more stringent access requirements for higher-value data and systems.

A key concept here is separation of duties, or SOD. The networking team should manage the network to achieve best performance while the security team manages access. When defined properly, SOD policies help ensure that no single person has complete control over a sensitive process or transaction. You don't want the network admin responsible for creating and managing user accounts to have permissions to delete security log entries, for example.



Pillar: Applications and workloads

Zero trust includes applications and workloads running on-premises, on mobile devices, and in the cloud. CISA recommends embedding security testing throughout the software development lifecycle and requiring automated testing of all applications.

Cloud native development practices help with a zero trust security model because these modular applications use a microservices architecture and containerization, which allows for isolation of components. Each service operates independently and has minimal access to other services unless explicitly granted.

More than **80%** of all data breaches today are attributed to employee error or negligence. By implementing zero trust, companies can help minimize the damage caused by careless end users.

Source: [SANS Institute](#)





Pillar: Data

Protecting data is a top reason why organizations adopt zero trust, so this pillar ties back into three focus areas by improving


- Availability, by helping make sure data is accessible when and where it's needed
- Confidentiality, by guarding against unauthorized access
- Integrity, by helping ensure data isn't altered, lost, or corrupted


In the CISA model, keys to success include enterprisewide automated data categorization and labeling, robust data loss prevention strategies that guard against theft of sensitive information, pervasive strong encryption, and continuous review of permissions.

Often, organizations already have many of these built into their databases—measures should include role-based access control; transparent encryption of data at rest and in transit; granular auditing to track and log user activity and changes to the database; and features to protect against common exploits, such as SQL injection attacks.

Finally, CISA's three cross-cutting capabilities—visibility and analytics, automation and orchestration, and governance—extend across all five pillars.

Visibility and analytics are about having the right data to inform policy decisions and respond quickly to problems. Automation and orchestration build on that data to handle security incidents as they arise. Governance capabilities help organizations manage and monitor their regulatory, legal, environmental, federal, and operational requirements to support risk-based decision-making.





Cultural realities of zero trust

CISA's governance layer is also about having the right people, processes, and technologies. Effective zero trust requires all three working in tandem, so if your current network security model grants a level of cross-asset mobility once users are inside the network, zero trust will likely be a significant change.

A sticking point is often new emphasis on identity and access management practices—think stricter password policies or even eliminating passwords altogether, multifactor authentication, and a more centralized approach to managing user identities and access privileges. People accustomed to looser controls may find these changes inconvenient, and, in fact, productivity may drop until users become accustomed to the new model.

Zero trust also involves more scrutiny of user activity and device health, which can raise privacy concerns among employees who feel their actions are being monitored too closely. Some will refuse to install mandated software on their personally owned devices.

And security, network operations, and application development pros aren't immune from resentment. You get the gist. It's a sea change where success hinges on user buy-in and strong executive leadership.

Strategies for a smooth transition include

Clear communication of the reasons behind adopting zero trust, emphasizing the business benefits. Openly address privacy concerns that employees might have and explain how zero trust protects their data. Consider trainings and info sessions on zero trust principles and best practices.

A phased rollout that gives employees, partners, and IT staff time to adjust. Prioritize implementing zero trust in a way that minimizes disruptions to workflows and maintains a positive user experience. Cloud-based technologies can help a lot here.

Defined separation of duties across the organization, with the goal of preventing users from accessing resources or performing actions outside their assigned roles. This can help prevent fraudulent activities across departments.

Leadership thanking people for their efforts. This can go a long way toward easing the transition. But executives also need to make zero trust a strategic initiative and allocate enough resources, work with IT teams to align controls with business objectives, and insist on key performance indicators to measure effectiveness.

6 indicators of zero trust success

In a recent [strategy guide](#), SANS Institute, a global provider of cybersecurity training and certification, provided six metrics to measure the effectiveness of your zero trust implementation.

- 1. Authentication success rate.** Measure the percentage of successful authentication attempts versus total attempts.
- 2. Policy compliance rate.** Measure the percentage of access requests that comply with your zero trust policies.
- 3. Number of lateral movement attempts.** Measure the number of detected attempts to move laterally within the network after initial access.
- 4. Time to detect and respond to incidents.** Measure the average time taken to detect and respond to security incidents.
- 5. User and entity behavior analytics anomalies.** Measure the number and severity of anomalies detected in user and entity behaviors to track potential security threats.
- 6. Stakeholder feedback.** Gather data from both technical and nontechnical roles across all groups, including cross-functional teams, end users, and managers.

6 best practices in architecting for zero trust

CISA points out that most large enterprises—including the federal government—face common challenges with legacy systems, where access and authorization are often based on fixed attributes and infrequently assessed. On a technology level, moving away from that outdated structure requires architectural changes.

Key best practices include

1

Timed and vetted access using an authorization framework

Say an employee logs on in the morning using an authorization service, such as OAuth, which issues tokens that are valid for specific systems and a limited period and that work in lieu of a password. When he needs to access a database, his entitlements for that system are confirmed by the token's authorization code.

2

Pervasive microsegmentation

The more granularly you can limit lateral movement without degrading performance, the better. CISA recommends distributed microperimeters and extensive microsegmentation based on how your applications are set up. This doesn't mean firewalls everywhere—techniques including virtual machines for each application, east/west traffic encryption, the creation of software-defined networks within the physical network, and intelligent routing algorithms can help isolate and secure segments without unduly slowing traffic flows.

3

Context-aware logging

AI can make a big difference here via adaptive authentication based on risk evaluation techniques that can help detect anomalies and, in response, step up authentication challenges. Ensure your providers can not only analyze logs in real-time to identify anomalies and potential threats but that they correlate events to spot complex patterns that could signal an advanced attack.

4

Pervasive encryption

Data is likely your most critical asset, and protecting data at rest, in transit, and in use demands end-to-end encryption paired with monitoring to detect unauthorized access attempts.

5

Least-privilege access

In the context of zero trust, least privilege is a core principle. Limiting access to only what's necessary for users, applications, and devices isn't about distrusting employees but about reducing risks. This approach minimizes the potential damage if a bad actor exploits stolen credentials, a compromised device, or a vulnerability.

6

A focus on device trustworthiness

Zero trust means no inherent trust for any device, even if it's inside the perimeter, company owned, or was previously granted access. Earning trust might involve meeting security posture requirements, such as having updated software, antivirus protection, and monitoring in place.


Got users who tend to lag on making software version or malware signature updates or resist installing security software on their personal devices? Zero trust will force their hands because an endpoint without the security profile defined by your policy simply won't be granted access. IT should manage endpoint security on company-owned devices, and compliance should be verified when new sessions are initiated.

Cloud providers and zero trust

Organizations that use cloud services may have a jumpstart in implementing zero trust. Hyperscalers offer preconfigured security solutions, templates, and tools, enabling faster deployment of zero trust controls and better scalability and centralized management.

In addition, many cloud providers have adopted GenAI-driven automation capabilities to minimize human errors and help security operations teams set up controls that take full advantage of native security features provided by cloud providers, such as firewalls, intrusion detection systems, encryption, continuous verification, least privilege access, identity lifecycle management, and microsegmentation.

A key enabling technology for cloud providers' zero trust efforts is [Zero Trust Packet Routing, or ZPR](#), an intent-based policy language that lets administrators define data access pathways for assets in the cloud. Traffic that isn't explicitly allowed by policy can't traverse the network.



ZPR benefits include

Human-readable policies written in natural language that are easy to understand, audit, and manage.

Decoupling of network security from network architecture, which helps reduce the risk of human error and configuration complexity.

Intent-based security, meaning ZPR uses security attributes to identify and organize resources and policies to control access to those resources.

Policy enforcement at the network level, regardless of network architecture changes or misconfigurations.



Other technologies that help with zero trust include zero trust network access, or ZTNA, and cloud access security brokers, or CASBs.

ZTNA, also known as software-defined perimeter, is a security approach that controls access to internal applications and resources in a much more granular way than a traditional VPN, which grants access to an entire network once a user is verified. ZTNA evaluates security credentials every time access to a resource is requested. The system considers context and may grant only partial access. If access is granted, it's via a secure session between the requesting entity and the specific asset. Then, activity and device health are continuously monitored for anomalous behavior that might indicate a threat.

A CASB uses machine learning to set a baseline of standard behavior for each user by monitoring typical activity over a period of time. The system then uses analytics to continuously compare behavior against the baseline to detect anomalous activity that may indicate a malicious insider or a compromised account. IT can set rules on when to suspend the account and when to raise an alert for follow-up.

Finally, cloud infrastructure entitlement management, or CIEM, focuses on managing identities and their permissions within cloud environments—IaaS, PaaS, and SaaS. Look for systems that integrate with identity standards, such as OAuth 2.0, security assertion markup language (SAML), and OpenID Connect, for authentication and authorization, and review what APIs your cloud providers offer to manage permissions and access control.

Bottom line, the zero trust “never trust, always verify” philosophy is fast becoming the norm. Organizations clinging to a perimeter mindset may lose the confidence of partners, customers, and employees. Fortunately, models and help are available, and cloud providers are leading the way.

Oracle Access Governance and zero trust

[Oracle Access Governance](#) enables a zero trust approach in the cloud by providing comprehensive identity governance with automation and fine-grained access control. It helps ensure that access is continuously evaluated and enforced dynamically based on context and policy, supporting the “never trust, always verify” principle.

[Oracle Advanced Authentication](#), [Oracle Adaptive Risk Management](#), and [Oracle Universal Authenticator](#) deliver robust capabilities to enforce a zero trust architecture. These solutions offer multifactor authentication (MFA), phishing-resistant MFA, adaptive authentication, and device-level MFA with passwordless options. Together, they provide dynamic security measures that adapt to user behavior, device integrity, and contextual risk, helping ensure secure access to critical resources.



How Oracle helps

With Oracle's security-first approach, explicit policies are required for access to Oracle Cloud Infrastructure (OCI). Each component is considered a resource within OCI, and access must be specifically granted. All communications within OCI are encrypted, and access rights are checked against existing policies. Those policies can be structured to grant extremely fine-grained access control for each resource, including implementing dynamic access.

OCI implements monitoring and auditing on cloud resources, letting you use existing object storage to conduct analysis, or you can employ your security information and event management tool of choice. Oracle Cloud Guard Instance Security provides automated responses to triggered events, helping speed reaction time to potential threats.

And the new OCI Zero Trust Landing Zone enables one-click provisioning for a secure, high-performing architecture for your cloud tenancy, with the deployment and hardened configuration of key services needed to meet certain requirements of zero trust.

Learn more

Connect with us

Call +1.800.ORACLE1 or visit oracle.com

Outside North America, find your local office at oracle.com/contact

Copyright © 2025 Oracle, Java, MySQL and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

