

Zero Trust in Practice

A phased roadmap for mid-size organisations, based on
NIST SP 800-207 and the CISA Zero Trust Maturity Model

SECURITY · NETWORK & COMMUNICATION

Published May 2025 · cybertechinsights.com

EXECUTIVE SUMMARY

The perimeter no longer describes how work happens

For decades, network security was built around a simple idea: everything inside the corporate network could be trusted, and everything outside could not. That model has been eroded from every direction. Staff work from home and on the move, applications run in several clouds and SaaS platforms, partners and contractors need access to specific systems, and attackers who steal a single set of credentials can often move freely once they are "inside".

Zero trust replaces location-based trust with explicit, continuous decisions about every access request. It is not a product that can be bought and switched on; it is an architecture and an operating approach that is adopted step by step. This whitepaper translates two authoritative public references — **NIST Special Publication 800-207, *Zero Trust Architecture*** (2020) and the **CISA Zero Trust Maturity Model, version 2.0** (2023) — into a practical, phased roadmap that a mid-size organisation can follow with a modest security team and an existing technology estate.

Key takeaways

- Zero trust removes implicit trust based on network location and decides each access request on its merits.
- Identity is the usual starting point: strong, phishing-resistant authentication and clean account lifecycles.
- Device health, segmented networks, application-level access and data protection follow in deliberate phases.
- Visibility, automation and governance cut across every pillar and determine how far maturity can progress.
- Progress is measured in maturity stages per pillar, not as a single "done" milestone.

CONCEPTS

What zero trust actually means

NIST SP 800-207 describes zero trust as a set of concepts designed to minimise uncertainty in enforcing accurate, least-privilege, per-request access decisions, in the face of a network that is assumed to be compromised. Its seven tenets can be summarised in plain terms:

- All data sources and computing services — including SaaS, cloud workloads and personal devices that access enterprise resources — are treated as resources to be protected.
- All communication is secured regardless of network location; being on the office network confers no privilege.
- Access to individual resources is granted on a per-session basis.
- Access is decided by dynamic policy that considers the identity of the user, the application or service, the state of the requesting device and other behavioural and environmental attributes.
- The organisation monitors and measures the integrity and security posture of all owned and associated assets.
- Authentication and authorisation are dynamic and strictly enforced before access is allowed.

- The organisation collects as much information as it can about the current state of assets, networks and communications, and uses it to improve its security posture.

The standard also describes the logical components that make these decisions. Understanding them helps buyers map existing tools to the architecture before purchasing anything new.

Component	Role in the architecture	Typical real-world examples
Policy engine (PE)	Makes the decision to grant, deny or revoke access to a resource, using policy and inputs such as identity, device state and threat intelligence.	Conditional access logic in an identity provider; the decision layer of a zero trust network access service.
Policy administrator (PA)	Acts on the decision: establishes or shuts down the communication path, issues session credentials or tokens.	Token issuance and session management in identity platforms; access brokers.
Policy enforcement point (PEP)	Enables, monitors and terminates the connection between the subject and the resource.	Access proxies, gateways, agents on endpoints, microsegmentation controls.
Data sources	Feed the decision: identity stores, device management, logs, threat intelligence, data classification.	Directory, endpoint management, EDR, SIEM, asset inventory.

FRAMEWORK

The CISA Zero Trust Maturity Model

NIST SP 800-207 explains *what* a zero trust architecture is. The CISA Zero Trust Maturity Model (ZTMM) helps organisations plan *how far* they have progressed. Version 2.0, released in April 2023, organises the work into five pillars and three cross-cutting capabilities, each assessed across four maturity stages:

Traditional, Initial, Advanced and **Optimal**. Although written for US federal civilian agencies, the model is publicly available and widely used as a vendor-neutral planning reference by private-sector organisations.

Pillar	What it covers	Practical early steps for a mid-size organisation
Identity	Users and service accounts, authentication, identity stores, risk assessment and access management.	Consolidate on one identity provider; enforce phishing-resistant MFA for administrators and remote access; automate joiner-mover-leaver processes.
Devices	Asset inventory, device compliance and health, protection of endpoints, servers and mobile devices.	Build a reliable device inventory; require managed, compliant devices for sensitive applications; deploy endpoint detection and response.
Networks	Segmentation, traffic management, encryption of traffic and network resilience.	Separate user, server, management and guest networks; encrypt internal traffic; begin replacing broad VPN access with application-level access.
Applications and workloads	Access to applications, secure development, application security testing and workload protection.	Put internal applications behind identity-aware access; integrate single sign-on; add security testing to the release process.
Data	Data inventory, categorisation, availability, access determination and encryption.	Identify the most sensitive data sets and where they live; apply labels; restrict and log access; review encryption at rest.

Cross-cutting capabilities

The model identifies three capabilities that support every pillar. In practice they are where many programmes stall, because they require integration between tools and teams rather than a single deployment.

- **Visibility and analytics** — collecting and correlating logs and telemetry from identity, devices, networks, applications and data so that policy decisions and investigations are based on evidence.
- **Automation and orchestration** — using automated workflows to provision access, respond to risk signals and enforce policy consistently, instead of relying on manual tickets.
- **Governance** — defining, enforcing and reviewing the policies, standards and responsibilities that the technical controls implement.

Reading the maturity stages

Stage	What it typically looks like
Traditional	Manually configured lifecycles and static policies; pillars managed in silos; limited automation and visibility.
Initial	Some automation of attribute assignment and policy enforcement; early cross-pillar integration.
Advanced	Automated controls across pillars where possible; centralised visibility; access adjusted to risk.
Optimal	Fully automated, just-in-time access lifecycles; dynamic policies informed by continuous, cross-pillar monitoring.

Few organisations need to reach Optimal everywhere. A realistic goal for a mid-size organisation is to move the pillars that protect its most valuable assets to Initial and then Advanced, while keeping the others at a deliberate, documented level.

PRIORITISATION

Where mid-size organisations should start

Mid-size organisations rarely have the budget or staff for a multi-year transformation programme. The roadmap should therefore be driven by risk and by the controls already owned. A useful sequence is:

- **Start with identity.** Most intrusions involve compromised credentials. Strong authentication, conditional access policies and timely removal of leavers deliver the largest early reduction in risk.
- **Protect the crown jewels first.** Identify a short list of “protect surfaces” — critical applications, data stores and administrative systems — and apply zero trust controls to these before tackling everything else.
- **Use what you already license.** Many identity, endpoint and collaboration subscriptions already include conditional access, device compliance and data labelling features that are not switched on.
- **Replace broad network access gradually.** Move one application group at a time from VPN or flat network access to identity-aware, per-application access, and decommission old paths as you go.
- **Design for operations.** Every new policy creates exceptions, alerts and support calls. Agree who owns them before enforcement begins.

Common pitfalls

- Treating zero trust as a single product purchase rather than an architecture spanning several existing tools.
- Enforcing strict policies without a monitoring-only period, which leads to outages and rapid rollback.
- Ignoring service accounts, APIs and machine identities, which often hold broad and long-lived privileges.
- Leaving legacy VPN and firewall rules in place after new access paths are introduced.
- Measuring progress by tools deployed rather than by maturity stage and risk reduced for priority assets.

BUYER GUIDANCE

Questions to ask zero trust vendors

- Which NIST SP 800-207 components does the product provide — policy engine, policy administrator, enforcement point, or data source?
- Which CISA ZTMM pillars and functions does it support, and at which maturity stage?
- What signals can it use in access decisions (identity risk, device compliance, location, behaviour), and from which third-party sources?
- How does it handle legacy applications and protocols that cannot use modern authentication?
- Can policies run in report-only or audit mode before enforcement, and how are exceptions managed?
- What logs does it produce, in what format, and how easily do they integrate with our existing SIEM?
- How does licensing scale with users, devices, applications and data volume?
- What happens to access if the vendor's cloud service is unavailable?

ROADMAP

A phased zero trust roadmap

Phase	Focus	Outcome
1. Assess and plan	Inventory identities, devices, applications and sensitive data; agree protect surfaces; baseline each ZTMM pillar; secure executive sponsorship.	Shared view of current maturity and a prioritised backlog.
2. Identity foundations	Single identity provider, phishing-resistant MFA for privileged and remote access, conditional access in report-only then enforced mode, automated leaver process.	Most access decisions based on verified identity.
3. Devices and access	Device inventory and compliance checks; per-application access for priority systems; begin retiring broad VPN access.	Access depends on both user and device posture.
4. Segmentation and data	Segment server and management networks; label and restrict sensitive data; encrypt internal traffic.	Reduced lateral movement and data exposure.
5. Automate and optimise	Centralise telemetry, automate responses to risk signals, review policies and reassess maturity regularly.	Continuous improvement towards Advanced maturity for priority pillars.

Checklist: zero trust readiness

- Protect surfaces and their owners are documented.
- Every pillar has a current and target ZTMM maturity stage.
- Phishing-resistant MFA protects administrators and remote access.
- Device compliance is checked before access to sensitive applications.
- Broad network access paths have a decommissioning plan.
- Sensitive data is labelled, and access to it is logged.
- Identity, device and network logs are correlated in one place.
- Policies are reviewed on a regular schedule, with exceptions tracked.

About Cyber Tech Insights: a B2B publication providing research and resources for technology and business decision-makers. This document is for general information only and does not constitute professional advice. © 2025 Cyber Tech Insights.