TECHNICAL EBOOK

intellizu

Zero Trust in Practice

Moving beyond the perimeter

Identity, Device, Network, Application & Data pillars · Migration planning
ZTNA vs VPN · Micro-segmentation · Privileged access · mTLS

Table of Contents

CHAPTER 1

The Perimeter Is Already Gone

For decades, network security assumed a defensible perimeter: corporate systems lived inside the network, untrusted traffic stayed outside, and the firewall enforced the boundary. That model has been eroding for years. Cloud infrastructure, SaaS adoption, remote work, and bring-your-own-device programs have made the perimeter a fiction. Zero Trust is the architectural response to that reality.

Why VPN-Centric Security Fails in Hybrid Environments

A VPN grants network access. Once a user is on the VPN, they have implicit access to everything on the network segment they've joined. That implicit trust is the problem. A stolen credential, a compromised device, or a malicious insider who authenticates to the VPN has the same network access as a legitimate user.

In hybrid environments, this gets worse. Applications live in multiple clouds, on-premises data centers, and SaaS platforms. VPN traffic must backhaul to a central point before reaching cloud services, increasing latency and degrading user experience. Organizations add split tunneling to fix the performance problem, which further undermines the security model by letting traffic bypass inspection.

- VPN grants network access, not application access — the scope of trust is too broad.
- Lateral movement after a breach is easy when internal network segments are flat.
- Performance and reliability trade-offs lead to split tunneling, creating policy gaps.
- Cloud and SaaS applications don't fit the hub-and-spoke VPN model well.

What Zero Trust Actually Means

Zero Trust is not a product category. It is an architectural principle: no user, device, or network location is implicitly trusted. Every request for access must be authenticated, authorized, and validated against policy — regardless of where the request originates.

The NIST definition (SP 800-207) is precise: Zero Trust assumes no implicit trust is granted based on network location or asset ownership. Every request is treated as if it originates from an untrusted network. Access is granted per-session to a specific resource, at the minimum necessary privilege, and evaluated dynamically.

■ Marketing vs. Reality

Many security vendors label their products 'Zero Trust' because it sells. A firewall is not Zero Trust. A VPN with MFA is not Zero Trust. Zero Trust requires per-request authorization based on verified identity and device state — not just network controls.

Never Trust, Always Verify — Applied Practically

The 'never trust, always verify' principle translates to specific architectural decisions. Instead of granting network access and letting users reach resources, access is granted to specific resources based on verified identity, device health, and context at the time of the request. The network location — inside or outside the office, on VPN or not — is irrelevant to the access decision.

- Authentication is required for every resource, including internal ones.
- Authorization is evaluated per-request, not per-session or per-network-connection.
- Device health is an input to the access decision, not just user identity.
- Minimum necessary privilege means users get access to what they need, not to the subnet.
- All traffic is treated as potentially hostile regardless of source network.

The Scope of the Shift

Zero Trust is not a single technology deployment. It is a multi-year architectural evolution that touches identity, devices, networking, applications, and data. Organizations that approach it as a product purchase rather than a sustained engineering program consistently underdeliver on the security outcomes.

The good news is that Zero Trust can be implemented incrementally, and each increment provides real security value. Identity controls come first and deliver the most value earliest. Network controls follow. Application-layer controls are the most granular and come last.

CHAPTER 2

The Five Pillars

Zero Trust architecture is typically organized around five pillars: Identity, Device, Network, Application, and Data. Each pillar represents a domain where implicit trust must be replaced with explicit verification. Understanding what each pillar means operationally — not just conceptually — is the starting point for planning.

Identity

Identity is the foundation. In a Zero Trust architecture, identity is the primary control plane. Who is this user, what are they allowed to do, and does the current context match what we expect for this kind of access?

Identity includes human users but also service accounts, workloads, and devices. Every entity requesting access to a resource needs a verified identity. The identity pillar covers authentication strength, identity lifecycle management, privileged access, and service-to-service identity.

- Strong authentication (phishing-resistant MFA or passwordless) for all human users.
- Just-in-time and just-enough-access for privileged operations.
- Workload identity for services — not shared passwords or long-lived credentials.
- Identity governance: provisioning, de-provisioning, and access reviews.

Device

Device trust means the access decision considers the health of the device making the request. A valid credential on a compromised device is not trustworthy. Device posture signals — patch level, EDR status, disk encryption, certificate presence — become inputs to access policy.

- Managed devices enrolled in MDM with health attestation.
- Conditional access policies that deny or limit access from unhealthy devices.
- Certificate-based device authentication that proves device identity to the network.
- BYOD handling: separate policies for personal versus managed devices.

Network

The network pillar replaces implicit network trust with explicit segmentation and traffic inspection. Rather than trusting traffic because it's on the internal network, traffic between systems is authenticated and authorized at the network layer.

- Micro-segmentation replaces flat internal networks with policy-governed segments.
- East-west traffic (service-to-service) is authenticated and logged, not implicitly trusted.
- DNS and traffic inspection for visibility into what's actually happening on the network.
- Network access is granted to specific resources, not to broad segments.

Application

The application pillar controls access at the application layer, independent of network controls. Applications should require authentication for every request, validate that the caller is authorized to perform the requested action, and not assume trust based on network location.

- Application access via identity-aware proxies (ZTNA), not VPN network access.
- API authentication with short-lived tokens, not long-lived API keys.
- Service-to-service authentication using workload identity and mTLS.
- Authorization at the application layer based on identity attributes, not IP ranges.

Data

The data pillar is about protecting data regardless of where it lives or how it's accessed. In a Zero Trust model, data classification drives access policy — the sensitivity of the data determines the strength of controls required to access it.

- Data classification that maps sensitivity levels to access control requirements.
- Encryption at rest and in transit for sensitive data, enforced by policy.
- Data loss prevention controls that apply based on data classification.
- Access logging at the data layer, not just the network layer.

■ Where to Start

Most organizations have partially addressed identity and have some network controls. The data pillar is frequently the least mature. But identity always comes first — strong authentication and access controls are the prerequisite for everything else.

CHAPTER 3

Identity as the New Perimeter

In a Zero Trust architecture, identity replaces the network as the primary security boundary. If you know exactly who — or what — is making a request, and you can verify their current context and health, you can make a meaningful access decision. This chapter covers what mature identity controls look like in practice.

Strong Authentication for Human Users

The baseline for Zero Trust identity is phishing-resistant authentication. Passwords alone are insufficient, and standard TOTP or push MFA — while better than nothing — are still susceptible to real-time phishing attacks. FIDO2/WebAuthn hardware keys or platform authenticators (Touch ID, Face ID, Windows Hello) provide origin-bound authentication that cannot be phished.

For most organizations, the practical priority is: enforce phishing-resistant MFA for privileged accounts first, then expand to the general employee population. Passkeys are the long-term direction — combining authentication and the second factor into a single phishing-resistant credential — and should be evaluated for new deployments on managed device fleets.

- FIDO2/WebAuthn for admins, engineers, and executives at minimum.
- Conditional access policies that require higher assurance for sensitive operations.
- Continuous session evaluation — re-authentication on risk signals, not just at login.
- Passwordless as an aspirational target: eliminates password-based attack surface.

Just-in-Time Access and Least Privilege

Persistent privileged access is one of the highest-risk configurations in any environment. An admin account that is always privileged is a target for credential theft, insider misuse, and account compromise. Just-in-time (JIT) access replaces persistent privilege with time-limited, purpose-scoped elevation.

In practice, JIT access means: users have standard access by default, elevation requires an approval workflow or policy-based self-service, elevated access is granted for a specific time window (hours, not days), and all activity during elevated sessions is logged and reviewable.

■ Common Gap

Most organizations have privileged accounts that are permanently elevated and rarely used. These accounts accumulate over years and are frequently not reviewed. An audit of accounts with admin or root-equivalent access typically surfaces 2-4x more accounts than security teams expect.

Privileged Access Management

PAM solutions like CyberArk, BeyondTrust, and HashiCorp Vault provide the infrastructure for JIT access, credential vaulting, and session recording. The core capabilities are: credential vaulting (no humans know production passwords), session proxying (privileged sessions go through a controlled gateway), and JIT workflow (access is requested and approved on demand).

- Vault all shared and service account credentials — eliminate knowledge of production passwords.
- Proxy privileged sessions through a recording gateway for auditability.
- Require MFA at the PAM layer as well as the IdP layer.
- Integrate PAM with SIEM for alert on anomalous privileged session activity.
- Review privileged account inventory quarterly and remove stale accounts.

Service Identity and Workload Authentication

Zero Trust identity is not only about human users. Services, workloads, and automated processes also need verified identities. Service-to-service authentication based on shared passwords or long-lived API keys is the equivalent of password-only authentication for machines — one leaked credential can compromise the entire service mesh.

Modern workload identity uses short-lived, automatically rotated credentials issued by a trusted identity authority. Cloud platforms provide this natively: AWS IAM roles with instance metadata, GCP workload identity federation, Azure managed identities. For on-premises or multi-cloud workloads, SPIFFE/SPIRE provides a platform-agnostic workload identity standard.

- Replace static service account passwords with platform-managed identity (IAM roles, managed identities, workload identity federation).
- Use short-lived tokens with automatic rotation — credentials that expire in hours, not years.
- Scope service permissions to least privilege — a service that reads from one S3 bucket should not have access to all S3 buckets.
- SPIFFE/SPIRE for environments that span clouds or include on-premises workloads.

CHAPTER 4

Device Trust

A verified identity is necessary but not sufficient for a Zero Trust access decision. A valid credential on a compromised or unmanaged device is still a risk. Device trust incorporates the health and management state of the device into the access decision.

MDM and Endpoint Health Signals

Mobile Device Management (MDM) provides the foundation for device trust: enrollment gives you visibility into the device, the ability to enforce policy, and the health signals needed for conditional access decisions. Without MDM coverage, you have no reliable way to assess device posture at access time.

Health signals that matter for access decisions include: OS version and patch level, disk encryption status, screen lock enforcement, EDR agent presence and health, compliance with security baselines, and certificate presence for device identity.

- MDM enrollment for all managed devices — without enrollment, health signals are unavailable.
- Compliance policies that define what 'healthy' means: encrypted disk, current patches, EDR running.
- Real-time compliance evaluation — health status checked at access time, not just at enrollment.
- Integration with IdP for conditional access: MDM compliance status flows to access policy.

Certificate-Based Device Authentication

Device certificates provide cryptographic device identity. A certificate issued to a specific device by a trusted CA proves the device is known and managed — not just that someone on the device has a valid user credential. Certificate-based authentication is a prerequisite for many Zero Trust network access patterns.

MDM platforms can automatically provision and manage device certificates at enrollment. The certificate is stored in the device's TPM or secure enclave, making it non-exportable. Network access controls, VPN clients, and ZTNA solutions can require a valid device certificate as an additional factor for network or application access.

■ Implementation Note

Many organizations have an internal PKI that issues user certificates but not device certificates. Extending the PKI to issue device certificates through MDM is often a low-friction step with high security value. The certificate lifecycle is managed automatically — no user action required.

BYOD vs. Managed Device Policies

BYOD is operationally convenient and expected in many organizations. It is also the hardest scenario for device trust. Personal devices cannot be fully managed — MDM enrollment on a personal phone gives the organization visibility and some policy enforcement, but not the same assurance as a corporate-managed device.

- Separate access tiers for managed vs. unmanaged devices — BYOD gets access to a limited set of resources appropriate to its lower trust level.
- Application-level MDM (MAM) for BYOD: manage the corporate app container without managing the personal device.
- Certificate issuance for managed devices only — certificates become the differentiator between managed and unmanaged device access policies.
- Clear off-boarding for BYOD: remote wipe of the corporate container when employment ends.

Conditional Access Policies

Conditional access is where device trust, identity, and context are combined into access decisions. A conditional access policy might say: allow access to the finance application only when the user has passed MFA, the device is MDM-enrolled and compliant, and the access is from an expected geography. Failure to meet any condition results in denied access or step-up authentication.

- Define access tiers: what resources require what combination of identity assurance and device health.
- Use named locations and risk signals (impossible travel, unfamiliar device) to trigger step-up authentication.
- Test conditional access policies in report-only mode before enforcement to surface unintended blocks.
- Audit policy coverage — every sensitive application should be covered by at least one conditional access policy.

CHAPTER 5

Network Segmentation That Works

Traditional network security relies on perimeter firewalls and implicit trust within the internal network. Zero Trust replaces implicit internal trust with explicit segmentation and authenticated traffic. This is one of the most operationally complex aspects of Zero Trust to implement — and one of the most valuable for stopping lateral movement.

The Problem With Flat Networks

A flat network is one where systems on the same network segment can communicate with each other without restriction. Most internal corporate networks are flat or close to it. An attacker who compromises one system on a flat network can reach every other system on the segment with no additional controls to bypass.

This is how most breaches expand from initial access to widespread compromise. The attacker doesn't need to defeat additional controls — they can move laterally using the same credentials and protocols that legitimate systems use. Segmentation limits the blast radius by forcing attackers to defeat additional controls to move between segments.

■ Lateral Movement Reality

In most breach investigations, the initial compromise is not the critical moment. The critical moment is when the attacker moves laterally to a system with the access they actually want. Micro-segmentation makes each lateral movement step significantly harder.

Micro-Segmentation

Micro-segmentation divides the network into small, policy-controlled segments. Rather than one internal network or a handful of VLANs, each application or workload group sits in its own segment, with explicit rules governing what can communicate with what.

Implementation approaches range from host-based firewalls with centralized policy management, to software-defined networking (SDN) platforms, to cloud security groups and network policies. The right approach depends on the environment — cloud-native workloads use cloud-native segmentation; on-premises workloads may use SDN or host-based controls.

- Start with application inventory: you can't segment what you haven't mapped.
- Identify communication flows between systems — what needs to talk to what.

- Create segments aligned with application boundaries and data sensitivity.
- Default-deny between segments, with explicit allow rules for required communication.
- Monitor inter-segment traffic for anomalies — unexpected communication paths are a significant signal.

East-West Traffic Controls

East-west traffic refers to traffic between systems within a data center or cloud environment — as opposed to north-south traffic, which crosses the perimeter. Traditional security focuses on north-south; Zero Trust requires equal attention to east-west.

East-west controls mean that even if an attacker reaches a system inside your environment, they cannot freely communicate with other systems. Every connection is subject to policy. This requires both network-layer controls and, ideally, application-layer authentication between services.

- Map all existing east-west traffic before applying controls — blocking undocumented flows will break applications.
- Use flow logs (VPC Flow Logs, NSG logs) to understand actual traffic patterns before writing policy.
- Start with monitoring-only mode: log policy violations before enforcing them.
- Alert on new east-west communication patterns that don't match baseline.

Service Mesh and mTLS

In microservices environments, a service mesh provides mutual TLS (mTLS) between services automatically. mTLS means both sides of a connection authenticate each other — the client proves its identity to the server, and the server proves its identity to the client. This is workload-to-workload Zero Trust at the application network layer.

Istio, Linkerd, and Consul Connect are the primary service mesh options. They inject a sidecar proxy alongside each service that handles mTLS transparently — the application code doesn't change. The mesh provides mTLS, traffic encryption, service-to-service authorization policies, and traffic observability.

- mTLS between all services in the mesh — unauthenticated service connections become impossible.
- Authorization policies at the service layer: Service A is allowed to call Service B's read endpoints; Service C is not.
- Observability built in: traffic between services is logged and traceable.
- Certificate rotation is automatic — short-lived certs issued by the mesh control plane.

CHAPTER 6

Application Access Patterns

How users and services access applications is the most visible layer of a Zero Trust architecture. This chapter covers the practical access patterns — ZTNA, application proxies, API authentication, and service-to-service auth — that replace VPN-based network access with identity-aware, application-level controls.

ZTNA vs. VPN: An Honest Comparison

Zero Trust Network Access (ZTNA) provides application-level access rather than network-level access. A user authenticates to the ZTNA system and gets access to specific applications — not to the network segment those applications live on. The application is effectively invisible to the user until access is granted.

	VPN	ZTNA
Access scope	Network segment	Specific application
Trust model	Network location	Identity + device + context
Lateral movement	Easy after auth	Prevented by design
Performance	Backhauling to hub	Direct-to-app (cloud-delivered)
Legacy app support	High	Varies by implementation
User experience	VPN client, all-or-nothing	Per-app, seamless for end users

Internal App Proxies

Not every internal application needs a full ZTNA platform. For web-based internal applications, an identity-aware reverse proxy can add authentication and authorization without changing the application. The proxy sits in front of the app, handles OIDC/SAML authentication, enforces MFA and device policy, and forwards authenticated requests to the app.

- Cloudflare Access / Pomerium / BeyondCorp Enterprise: cloud-delivered proxies with identity-aware access for internal web apps.
- OAuth2 Proxy: open-source option that integrates with any OIDC provider.
- Nginx auth_request: flexible but requires maintaining a validation service.

- Works for any HTTP/HTTPS application without application code changes.

The proxy pattern is especially valuable for legacy internal applications that predate SSO integration. Adding a proxy in front is often faster and lower-risk than modifying the application itself.

API Gateway Authentication

APIs are a significant attack surface that often receives less attention than user-facing applications. API authentication in a Zero Trust model requires every API call to carry a verifiable credential — typically a short-lived JWT issued by the identity provider or a service identity token.

- OAuth 2.0 with short-lived bearer tokens: tokens expire in minutes to hours, not months.
- API gateway as enforcement point: validate tokens centrally before requests reach services.
- mTLS for service-to-service API calls in addition to token authentication — belt-and-suspenders for sensitive internal APIs.
- Eliminate long-lived API keys in favor of dynamic credentials — keys that don't expire are a liability.
- Log all API authentication failures — patterns of invalid tokens often indicate credential theft or misuse.

Service-to-Service Authentication Patterns

Services calling other services need authentication patterns that don't rely on humans. The options are: platform-managed identity (cloud IAM roles), service mesh mTLS, or explicit token exchange using a standards-based protocol like SPIFFE/SVID.

The wrong pattern — shared static credentials, IP-based trust, or trusting any traffic that arrives on the internal network — is still common. Each of these creates risk: a leaked credential grants unlimited service-to-service access, IP-based trust fails once any internal host is compromised, and implicit internal trust is exactly what Zero Trust is designed to eliminate.

■ Migration Path

Auditing service-to-service authentication almost always surfaces hardcoded credentials in configuration files or environment variables. Replace these with platform-managed identity where available, or rotate to short-lived credentials issued by a secrets manager with automatic rotation. Static secrets should be treated as a migration target, not a permanent state.

CHAPTER 7

Getting There From Here

Zero Trust is a destination, not a product you can install. Most organizations are starting from a perimeter-centric model with years of technical debt in identity, networking, and access management. This chapter is about the practical migration path: what to prioritize, what to defer, and how to measure progress.

The Sequencing Principle: Identity Always Comes First

Every Zero Trust framework — NIST, CISA, Google BeyondCorp — agrees on sequencing: identity controls come before network controls, and network controls come before application and data controls. The reason is practical: you cannot make meaningful access decisions without verified identity. Everything else builds on that foundation.

Starting with network micro-segmentation before fixing identity is a common mistake. Sophisticated network controls still fail if the identity feeding access decisions is untrustworthy — compromised credentials or unverified devices can still navigate well-segmented networks if identity is the only control.

■ Priority Order

1. Strong authentication for all users (especially privileged). 2. Conditional access policies that incorporate device health. 3. Privileged access management with JIT elevation. 4. Application access via identity-aware proxies or ZTNA. 5. Network micro-segmentation and east-west controls. 6. Service mesh mTLS and workload identity.

Assessment: Knowing Where You Are

Before prioritizing changes, you need an honest inventory of where implicit trust currently lives in your environment. This is more work than it sounds — most organizations have a gap between documented architecture and actual access patterns.

- Authentication path inventory: how do users authenticate to each system, and what bypasses the IdP entirely?
- Privileged account inventory: every account with elevated permissions, when last used, whether access is still needed.
- Service account audit: every service-to-service credential, its scope, and how it's stored.

- Network topology: current segmentation boundaries, what can communicate with what without restriction.
- Application access patterns: which apps require VPN, which are directly internet-accessible, which have no authentication controls.

What to Prioritize First

The highest-value, lowest-disruption starting points for most organizations are:

- Phishing-resistant MFA for all privileged accounts. Small population, high value, low disruption.
- Conditional access policies that block access from non-compliant devices for sensitive applications.
- Privileged access management: vault production credentials and implement JIT elevation for admin tasks.
- MFA enforcement for all employees — broad population, meaningful reduction in credential-based breach risk.
- ZTNA for remote access to internal applications — replaces VPN for the most common access pattern.

What to Defer

Not everything needs to happen in year one. Deferring the right items lets you deliver early value while managing operational risk.

- Full micro-segmentation of the entire internal network — start with the most sensitive systems and expand.
- Service mesh deployment across all services — prioritize services that handle sensitive data first.
- Replacing legacy applications with modern auth-aware alternatives — proxy patterns can bridge the gap while replacement is planned.
- Data classification and DLP — valuable but dependent on identity and access controls being stable first.

Measurement: How to Know It's Working

Zero Trust programs need metrics to demonstrate progress and identify gaps. The right metrics measure actual security outcomes, not just deployment counts.

- MFA coverage rate: percentage of authentication events going through strong MFA. Gaps reveal bypass paths.
- Device compliance rate: percentage of access events from compliant, managed devices.

- Privileged access with JIT: percentage of privileged operations performed through JIT workflows versus persistent elevated access.
- Micro-segmentation coverage: percentage of workloads in segmented environments with explicit inter-segment policies.
- Lateral movement simulation: periodic red team exercises testing whether an attacker with one compromised system can reach sensitive resources.
- Mean time to detect (MTTD) for anomalous access: Zero Trust generates more authentication and authorization data — use it.

■ Long-Term Horizon

Zero Trust is not a project with a completion date. It is an ongoing program that evolves with your environment. The goal is continuous improvement toward less implicit trust, measured by the metrics above, with a regular review cycle that identifies where the current gaps are and what to address next.

How Intellizu Can Help

Moving to Zero Trust surfaces questions that go beyond policy documents: where does implicit trust actually live in your environment, which service-to-service connections have never been inventoried, what does your real authentication and access landscape look like versus how it's documented. These questions require a structured look at how identity, access, and network controls actually work — not just how they're supposed to work.

Intellizu's Systems Assessment is a focused engagement that maps your current auth and access landscape across identity, devices, network, and application layers. It identifies where implicit trust exists today, what the highest-value controls to add first are, and produces a prioritized roadmap your team can act on. For teams that need ongoing engineering capacity — implementing ZTNA, configuring conditional access policies, deploying mTLS in a service mesh, or running a phased migration — we work as an engineering retainer: embedded support that brings implementation alongside strategy.

If you're planning a Zero Trust program and want a second opinion on your architecture, or a partner to help execute it, we'd be glad to talk.

intellizu.com