WHITE PAPER

Zero Trust in a Post-Perimeter World: Securing Hybrid & SaaS Environments

Why identity is the new firewall - and how to build security architecture fit for the cloud-first enterprise

A Digital Resilience Strategic Advisory Report

Digital Resilience | www.digitalresilience.co.za | April 2025

Target Audience: IT Managers, Cloud Security Engineers, Security Directors

EXECUTIVE SUMMARY

The network perimeter is gone. Not weakened, not porous - structurally dissolved. The enterprise of 2025 runs its workloads across three or four cloud providers, relies on over 100 SaaS applications, and employs a workforce that connects from home networks, airports, and client offices using personal and managed devices alike. The castle-and-moat security model that defined enterprise defence for three decades was not designed for this world. It cannot be adapted to it. It must be replaced.

Zero Trust Architecture (ZTA) is that replacement - not a product, not a vendor category, but an architectural philosophy grounded in a single, radical premise: **no user, device, or workload is trusted by default, regardless of where it sits relative to the network boundary.** Identity becomes the new perimeter. Verification becomes continuous. Least privilege becomes the enforced norm rather than the aspirational one.

This white paper sets out the case for Zero Trust with urgency and specificity. It examines why traditional perimeter defences fail in hybrid and SaaS environments, provides a practical framework for implementing ZTA across multi-cloud infrastructure, addresses the critically underexamined threat posed by non-human identities, and presents a model for continuous verification that balances security rigour with operational usability. It concludes with a sequenced set of recommendations grounded in what leading security organisations are doing today.

100+ SaaS applications used by the average enterprise (Okta Business at Work, 2024)	**80%+** of breaches involve compromised credentials (Verizon DBIR, 2024)
10:1 ratio of non-human to human identities in a modern enterprise environment	**USD 4.5M** average cost saving when a mature Zero Trust programme is in place (IBM, 2024)

> *"Zero Trust is not a technology purchase. It is a commitment to a fundamentally different relationship between your organisation and the concept of trust itself."*
>
> **- Yaron Assabi**

THE DEATH OF THE PERIMETER – WHY THE OLD MODEL FAILS

The traditional network security model was built on a clear and defensible assumption: everything inside the network boundary can be trusted; everything outside it cannot. Security investment focused on making that boundary as strong as possible - firewalls, DMZs, intrusion detection systems, and VPN gateways that extended the trusted zone to remote workers.

That assumption no longer holds. The 'trusted inside' is a fiction. When critical data lives in Salesforce, Microsoft 365, and a dozen other SaaS platforms, there is no meaningful 'inside' to protect. When workloads run across AWS, Azure, and GCP simultaneously, the network boundary is not a line but a scattered constellation of endpoints, APIs, and services - none of which can be definitively declared trusted or untrusted based on network location alone.

The VPN Problem

Many organisations responded to the collapse of the perimeter by extending it via VPN - treating remote access as a matter of pulling employees back inside the trusted zone before they work. This approach is inadequate for two structural reasons. First, it grants excessive trust to anyone who authenticates successfully: a compromised VPN credential gives an attacker the same broad access as a legitimate employee. Second, it does not address the reality that the most sensitive data may not live inside the VPN tunnel at all - it lives in SaaS applications that the VPN does not protect.

> **Case Study: Colonial Pipeline (2021)**
> The Colonial Pipeline ransomware attack - which caused fuel shortages across the US East Coast and resulted in a USD 4.4 million ransom payment - originated from a single compromised VPN password. The account had no multi-factor authentication enabled. Once authenticated, the attacker moved laterally through the network without restriction. The perimeter had been breached through its own front door. A Zero Trust architecture requiring continuous verification and least-privilege access would have dramatically contained the blast radius.

The Implicit Trust Problem

The deeper issue is implicit trust - the assumption baked into traditional network architectures that authenticated users and devices retain their trustworthiness for the duration of a session. An employee who authenticates at 9:00am is trusted at 11:00am, even if their credentials were stolen at 10:30am. A device that passes a compliance check at login remains 'trusted' even if its security posture degrades during the day.

Credential theft has become the dominant initial access vector precisely because it exploits this implicit trust. According to Verizon's 2024 Data Breach Investigations Report, over 80% of breaches involving hacking used stolen credentials. Attackers do not need to break through the perimeter - they walk through it with a valid key.

Regulatory and Insurance Pressure

The regulatory landscape is accelerating ZTA adoption. The EU's NIS2 Directive (effective October 2024) requires organisations in critical sectors to implement access control measures proportionate to their risk exposure - language that increasingly maps to Zero Trust principles in regulatory guidance. The EU DORA framework adds specific requirements for financial entities around identity and access management. In the US, the Biden Administration's 2021 Executive Order on Improving the Nation's Cybersecurity explicitly mandates ZTA adoption across federal agencies, with commercial sector follow-on expected.

Cyber insurance underwriters are moving in parallel. Leading insurers now require multi-factor authentication, privileged access management, and network segmentation as baseline conditions for coverage - all components of a ZTA approach. Organisations that cannot demonstrate these capabilities face either premium increases, reduced coverage limits, or outright exclusions.

IMPLEMENTING ZTA IN DECENTRALISED, MULTI-CLOUD, AND SAAS ENVIRONMENTS

NIST Special Publication 800-207 defines Zero Trust Architecture around five core pillars: Identity, Devices, Networks, Applications and Workloads, and Data. Each pillar represents a control plane that must be addressed. In a modern hybrid environment, implementing these pillars requires both a coherent philosophical approach and provider-specific technical execution.

The Five Pillars of Zero Trust

Pillar	Core Principle	Key Controls
Identity	Every user and workload is continuously authenticated and authorised	MFA, Conditional Access, Identity Governance, PAM
Devices	Device health and compliance are verified before granting access	MDM/UEM, EDR integration, device posture assessment

Pillar	Core Principle	Key Controls
Networks	Micro-segmentation reduces blast radius; no implicit east-west trust	Service mesh, SASE, ZTNA, network segmentation policy
Applications	Access granted at the application layer, not the network layer	CASB, SSE, application-layer proxy, API gateway controls
Data	Data classification drives access decisions wherever data lives	DLP, data classification, encryption at rest and in transit

Multi-Cloud Implementation Challenges

Each major cloud provider - AWS, Azure, GCP - has its own identity and access management primitives, policy languages, and enforcement mechanisms. Achieving consistent Zero Trust policy across providers requires an abstraction layer that normalises these differences without sacrificing granularity.

Cloud Infrastructure Entitlement Management (CIEM) platforms address one of the most pervasive multi-cloud risks: over-permissioned identities. In cloud environments, identities - both human and non-human - routinely accumulate permissions far beyond what their roles require, because provisioning is easy and de-provisioning is rarely prioritised. CIEM tools continuously analyse entitlements and recommend right-sizing to enforce least privilege at cloud scale.

Federated identity is the practical foundation of multi-cloud Zero Trust: using a single authoritative Identity Provider (Okta, Microsoft Entra ID, Ping Identity) as the source of truth across all cloud environments. Service mesh technologies such as Istio extend this principle to workload-to-workload communication, ensuring that microservices authenticate each other using cryptographic certificates rather than implicit network trust.

SaaS-Specific Challenges and Controls

SaaS environments present a distinct challenge: the organisation does not control the application infrastructure. Zero Trust in SaaS requires working at the access layer rather than the application layer, using Cloud Access Security Brokers (CASB) and Security Service Edge (SSE) platforms to enforce policy on the connection between users and SaaS applications.

Shadow SaaS - the adoption of unsanctioned applications by employees outside IT governance - is a systematic gap in ZTA coverage. A user who accesses company data through an unapproved file-sharing application bypasses every Zero Trust control applied to sanctioned SaaS. Discovery and governance of shadow SaaS through CASB-based visibility tools is therefore a prerequisite for effective ZTA, not an optional enhancement.

OAuth and the Over-Permission Problem
OAuth grants - the permissions employees extend to third-party applications when they click 'Sign in with Google' or 'Connect to Microsoft 365' - represent a significant and systematically undermanaged attack surface. The average enterprise has thousands of OAuth grants outstanding, many granting broad access to email, calendar, and file storage. A compromised third-party application inherits all of those permissions. Regular OAuth grant audits and governance workflows should be standard practice in any mature Zero Trust programme.

A Phased ZTA Implementation Roadmap

Phase	Timeline	Focus	Key Deliverables
1 - Foundation	0-3 months	Identity consolidation	Universal MFA, single IdP, privileged access inventory
2 - Device Trust	3-6 months	Device compliance enforcement	MDM enrolment, EDR integration with access decisions, device posture baseline
3 - Network	6-12 months	Micro-segmentation	ZTNA deployment, east-west traffic controls, service mesh pilot
4 - Application	12-18 months	App-layer access controls	CASB/SSE deployment, OAuth governance, shadow SaaS discovery
5 - Data	18-24 months	Data-centric policy	Data classification schema, DLP enforcement, encryption policy review

IDENTITY-BASED ATTACKS AND THE NON-HUMAN IDENTITY CRISIS

If the perimeter is dead, identity is the battlefield. Threat actors have adapted accordingly: credential-based attacks now account for the majority of initial access vectors, and the sophistication of those attacks has increased dramatically. Understanding the full scope of the identity attack surface - including the vast and underprotected population of non-human identities - is essential for any organisation building a Zero Trust programme.

The Modern Credential Attack Playbook

Traditional credential attacks relied on brute force or phishing to obtain usernames and passwords. The modern playbook is considerably more sophisticated, targeting the authentication infrastructure itself rather than individual credentials.

Adversary-in-the-Middle (AiTM) phishing uses reverse proxy frameworks (Evilginx, Modlishka) to intercept authentication flows in real time, capturing session tokens after MFA has been completed. The attacker does not steal the password or the MFA code - they steal the authenticated session, making traditional MFA completely ineffective against this vector.

MFA fatigue attacks exploit push-based MFA by bombarding a user with approval requests until, through fatigue or confusion, they approve one. This technique was used in the 2022 Uber breach, where an attacker obtained credentials through social engineering and then overwhelmed the target with MFA push notifications until one was approved.

The response is not to abandon MFA but to upgrade it. **FIDO2/WebAuthn and passkeys** are phishing-resistant by design: authentication is bound to the specific origin and cannot be intercepted by a proxy, and there is no shared secret that can be phished. For high-risk accounts and applications, hardware security keys (YubiKey, Google Titan) represent the current gold standard.

> *"Attackers are not breaking your door down. They are walking through it with credentials that look completely legitimate. The only defence is to never assume those credentials are still legitimate."*
>
> **- Yaron Assabi**

The Non-Human Identity Explosion

The identity problem in most enterprises extends far beyond human users. Service accounts, API keys, OAuth tokens, CI/CD pipeline credentials, cloud IAM roles, and Kubernetes service accounts - collectively termed non-human identities - outnumber human identities by approximately 10:1 in a modern enterprise. They are disproportionately powerful and disproportionately ungoverned.

Non-human identities are typically created for convenience, provisioned with broad permissions to avoid integration failures, and then forgotten. They rarely appear in identity governance reviews. Their credentials are often static and long-lived. When they are compromised - most commonly through leaked API keys in public GitHub repositories, which occur thousands of times daily - the blast radius can be severe.

10:1 Non-human to human identity ratio in the average enterprise	**6 hrs** Average time for a leaked secret on GitHub to be exploited (GitGuardian, 2024)
72% of security teams report non-human identities as their greatest IAM risk	**USD 4.2M** Average breach cost when non-human credentials were involved (IBM, 2024)

Non-Human Identity Governance

Addressing the non-human identity attack surface requires treating machine identities with the same governance rigour as human identities - which means inventory, least privilege, rotation, and continuous monitoring:

- **Secrets management platforms** (HashiCorp Vault, AWS Secrets Manager, Azure Key Vault) centralise the storage and lifecycle management of API keys, passwords, and certificates - enabling automatic rotation and eliminating hardcoded credentials in code.
- **Workload Identity Federation** allows cloud workloads to authenticate using short-lived, automatically rotated tokens rather than long-lived service account keys - eliminating a major class of credential theft risk.
- **Continuous secrets scanning** in code repositories and CI/CD pipelines detects credential exposure within seconds of a commit, enabling rapid remediation before leaked secrets can be exploited.
- **Non-human identity lifecycle management** - provisioning, review, and de-provisioning workflows for service accounts that mirror those applied to human identities - should be a standard component of any IAM programme.

CONTINUOUS VERIFICATION VS. THE TRUST-ONCE MODEL

The philosophical heart of Zero Trust is the rejection of the trust-once model - the assumption that a user who authenticates successfully at the beginning of a session remains trustworthy for its duration. Continuous verification replaces this assumption with an ongoing, risk-proportionate assessment of whether access should be maintained.

Why Session-Based Trust Fails

Session tokens are the new passwords. Once a user authenticates and receives a session token - the cryptographic credential that tells an application 'this user has already authenticated' - that token represents trusted access for its entire lifetime. Many SaaS applications issue tokens valid for 24 hours, 7 days, or longer. An attacker who steals a session token through pass-the-cookie attacks, AiTM phishing, or browser-based malware inherits full, authenticated access for that entire window.

The 'assumed breach' mentality of Zero Trust treats this risk as a design constraint rather than an edge case. If it must be assumed that credentials or sessions may already be compromised, then security cannot rely on the validity of authentication as a one-time event. Trust must be earned continuously, not granted at login and forgotten.

Risk-Adaptive Authentication in Practice

Continuous verification does not mean asking users to re-authenticate every few minutes - which would create intolerable friction and drive workarounds. It means continuously evaluating contextual signals and adjusting trust levels accordingly, requiring additional verification only when the risk profile changes materially.

Risk signals used in adaptive authentication include: device health and compliance status, geographic location and velocity (impossible travel), time of day and access pattern anomalies, data sensitivity level of the resource being accessed, user behaviour analytics (unusual typing patterns, data volume, access sequences), and threat intelligence signals (known malicious IPs, Tor exit nodes).

> **Step-Up Authentication: The Proportionate Response**
> A mature continuous verification architecture implements step-up authentication: a mechanism that allows low-risk operations to proceed with minimal friction while requiring additional verification for higher-risk actions mid-session. A user reading internal documentation on a managed device needs no additional friction. The same user attempting to export 50,000 customer records, access payroll data from an unrecognised device, or modify admin credentials triggers an immediate step-up request - additional MFA, re-authentication, or manager approval - before the action is permitted.

The User Experience Imperative

Security programmes that ignore user experience consistently fail. When friction is high enough, users find workarounds - sharing credentials, using personal devices, routing around CASB controls through personal email. These workarounds create exactly the security gaps that Zero Trust is designed to close. Designing for low-friction continuous verification is therefore not a nice-to-have - it is a prerequisite for the programme's effectiveness.

The ideal state is transparent, invisible continuous verification: contextual signals are evaluated in the background and trust is maintained without user interaction in the vast majority of sessions. Step-up authentication fires only when the risk profile genuinely warrants it, making those friction events meaningful rather than routine. Single Sign-On (SSO) reduces authentication events dramatically while centralising control, and passwordless authentication (passkeys, hardware keys) eliminates the most common point of failure entirely.

Measuring Zero Trust Maturity

Maturity Level	Identity & Access	Device Trust	Network	Data
1 - Initial	Password-only, no MFA	No device checks	Flat network, VPN	No classification
2 - Developing	MFA for some apps, AD-centric	Basic MDM	Some segmentation	Manual tagging
3 - Defined	Universal MFA, IdP established	Compliance-gated access	ZTNA deployed	Automated classification
4 - Managed	Adaptive auth, CIEM, PAM	Continuous posture	Micro-segmentation	DLP enforced
5 - Optimising	Passkeys, FIDO2, AI-driven auth	Zero-trust device graph	Zero-trust by default	Data-centric policy

STRATEGIC RECOMMENDATIONS FOR SECURITY LEADERS

Zero Trust is a multi-year programme, not a product deployment. The following recommendations are sequenced for organisations at varying maturity levels, with an emphasis on high-impact, low-regret actions that build foundational capability before more advanced controls.

Immediate Actions (0-3 Months)

- **Enforce MFA universally - no exceptions.** Prioritise internet-facing applications, administrator accounts, and remote access systems. Where push-based MFA is deployed, evaluate upgrade paths to number-matching or FIDO2 to counter MFA fatigue attacks.

- **Conduct an identity audit.** Inventory all identities - human and non-human - with access to production systems. Identify accounts with excessive privileges, stale accounts, and service accounts with long-lived credentials. The findings will almost certainly be surprising.
- **Eliminate shared credentials.** Shared service accounts and generic admin passwords are Zero Trust violations by definition. Each identity - human or machine - must be individual, auditable, and accountable.
- **Begin OAuth grant governance.** Use your IdP or CASB tooling to inventory all third-party OAuth grants across Microsoft 365 and Google Workspace. Revoke any grants that are excessive, stale, or from unrecognised applications.

Medium-Term Investments (3-12 Months)

- **Deploy a ZTNA solution** to replace or complement your VPN. ZTNA grants access to specific applications rather than broad network segments, dramatically reducing the blast radius of compromised credentials. Modern SASE platforms bundle ZTNA with CASB and SWG in a single cloud-delivered service.
- **Implement adaptive, risk-based access policies** in your IdP. Use device compliance, location, and risk signals to dynamically adjust authentication requirements. High-risk contexts should trigger step-up authentication; low-risk contexts should be frictionless.
- **Deploy a secrets management platform** and begin migrating hardcoded credentials from application code and configuration files. Integrate secrets scanning into your CI/CD pipeline immediately to prevent new exposures.
- **Pilot micro-segmentation** for your highest-value application environments. Start with one business-critical system and demonstrate the blast-radius reduction before scaling across the estate.

Strategic Positioning (12+ Months)

Organisations that establish the identity, device, and network foundations of Zero Trust in the near term will be positioned for the more ambitious controls that define mature ZTA programmes: data-centric policy, workload identity federation at scale, and AI-driven continuous risk scoring. The investment pays forward: every component of Zero Trust infrastructure builds on the last, making the programme progressively more capable and more cost-effective over time.

The board-level framing matters. Zero Trust is not a security cost - it is a business enablement investment. Organisations with mature ZTA programmes can safely extend access to partners, contractors, and acquired entities without the months-long network integration projects that traditional perimeter models require. They can adopt new SaaS applications without security reviews that delay business value by quarters. They can demonstrate to insurers, regulators, and customers a security posture that is measurable, auditable, and continuously improving.

SOURCES AND REFERENCES

1. **NIST**. *Special Publication 800-207: Zero Trust Architecture*. https://csrc.nist.gov/publications/detail/sp/800-207/final

2. **Verizon**. *Data Breach Investigations Report 2024*. https://www.verizon.com/business/resources/reports/dbir/

3. **Okta**. *Business at Work 2024: SaaS Application Trends*. https://www.okta.com/businesses-at-work/

4. **IBM Security**. *Cost of a Data Breach Report 2024*. https://www.ibm.com/reports/data-breach

5. **GitGuardian**. *State of Secrets Sprawl 2024*. https://www.gitguardian.com/state-of-secrets-sprawl

6. **CISA**. *Zero Trust Maturity Model Version 2.0*. https://www.cisa.gov/zero-trust-maturity-model

7. **Microsoft**. *Entra ID and Zero Trust - Technical Reference*. https://learn.microsoft.com/en-us/security/zero-trust/

8. **Gartner**. *Market Guide for Zero Trust Network Access (ZTNA), 2024*. https://www.gartner.com/en/documents/zero-trust

9. **European Commission**. *NIS2 Directive - Official Text*. https://digital-strategy.ec.europa.eu/en/policies/nis2-directive

10. **FIDO Alliance**. *FIDO2 and WebAuthn Technical Specification*. https://fidoalliance.org/fido2/

Partner with Digital Resilience

Digital Resilience helps organisations design and implement Zero Trust architectures that are pragmatic, phased, and fit for their specific hybrid and SaaS environments. Our advisory services span ZTA strategy and roadmap development, identity governance programme design, non-human identity risk assessment, ZTNA and SASE vendor selection, and continuous verification architecture review.

Whether you are beginning your Zero Trust journey, accelerating an existing programme, or seeking an independent assessment of your current architecture's resilience, our practitioners bring the technical depth and strategic clarity to move your programme forward with confidence.

Protecting your Business Ecosystem

yaron@digitalresilience.co.za | www.digitalresilience.co.za