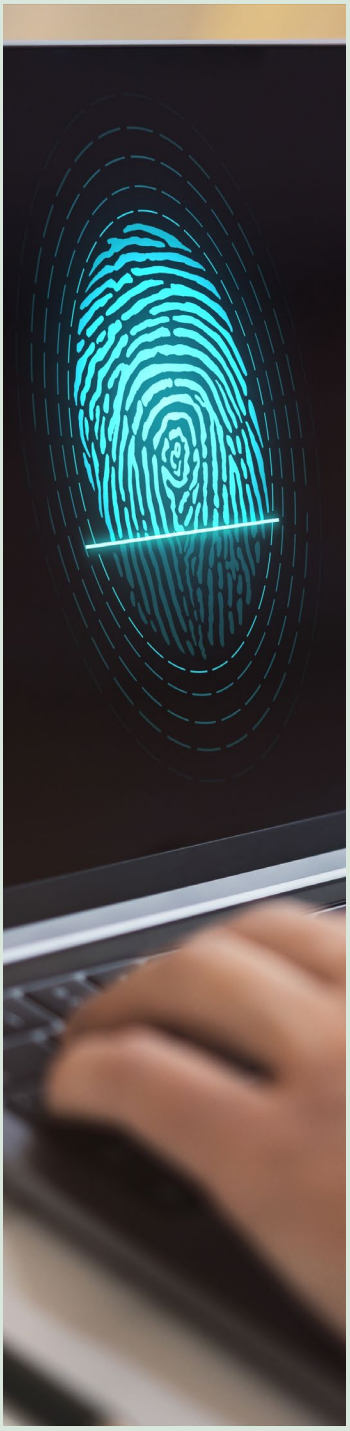
WHITE PAPER



IDENTITY SECURITY IN THE AI ERA: FROM STATIC ACCESS TO CONTINUOUS CONTROL

Abstract

Enterprise environments no longer operate within clearly defined network boundaries. Identity systems now determine how users, services, and automated processes interact across distributed platforms.[1]

As identity becomes the primary control plane for access, attack patterns have evolved. Rather than relying solely on infrastructure exploitation, attackers increasingly gain access through legitimate identity pathways such as credential compromise, session misuse, and impersonation. Credential abuse was the initial access vector in 22% of confirmed breaches, highlighting identity compromise as one of the leading pathways into enterprise environments.[2]

Artificial intelligence increasingly shapes how access is created, exercised, and abused. Access is now exercised continuously by systems, integrations, and automated workflows operating beyond direct human oversight. Attackers are exploiting these conditions, using automation and AI-enabled techniques to accelerate impersonation, credential misuse, and privilege escalation.[3]

Most identity models were built around assumptions that access is requested deliberately, used only when required, and reviewed before it becomes risky. Those assumptions no longer hold. Access is still assigned using static rules, but permissions often persist long after roles, systems, and operational context have changed. Over time, this creates large volumes of technically valid access that are misaligned with current needs and increasingly opaque to traditional control mechanisms.[4]

The result is a widening gap between how identity is governed and how access operates across enterprise environments. Identity governance can no longer rely on static access assumptions when access relationships continuously evolve across users, services, APIs, and automated workloads.

This white paper outlines an identity security operating model for the AI era and explains how organisations can continuously align identity, access, and behaviour across evolving enterprise environments.

themissinglink®
An Infosys company

Infosys
Navigate your next

Industry Problem

1. The Perimeter Has Moved

Security models have traditionally relied on network-based boundaries, where systems were segmented and users inside those boundaries were implicitly trusted. As applications and data extended into cloud and distributed environments, this model became increasingly ineffective.[1]

Access is now controlled through identity systems that coordinate authentication, authorisation, and policy enforcement across environments. Authentication establishes who or what is requesting access, authorisation defines permitted actions, and policy enforcement determines how access is applied.

Modern enterprise access increasingly relies on federated identity protocols such as SAML, OAuth 2.0, and OpenID Connect (OIDC), extending authentication and authorisation decisions across SaaS platforms, cloud services, APIs, and distributed workloads. As organisations adopt more API-driven and cloud-native operating models, identity systems increasingly govern token-based access between users, applications, services, and automated processes operating across multiple trust boundaries.

Because these controls operate consistently across environments, identity systems define the effective boundary for access. This aligns with Zero Trust principles, where access decisions are based on identity and context rather than network location.[1]

Identity now sits at the centre of security, operational resilience, and regulatory compliance across modern enterprise environments. These controls enable traceability, enforcement, and auditability across distributed access environments, making identity governance foundational to both security and operational assurance. Failures in identity governance therefore expose organisations not only to breaches, but also to regulatory and operational impact.[4]

Breach data reflects this change. Most incidents now involve the use of valid credentials rather than exploitation of technical vulnerabilities, indicating that the dominant failure point has moved from infrastructure to identity control.[2]

2. Identity Systems Were Built for a Different Operating Model

Traditional identity frameworks were built around predictable patterns of human behaviour. Access was intermittent, aligned to stable roles, and constrained by session-based interaction. These assumptions shaped how identity systems grant, evaluate, and revoke access.[1]

Those assumptions no longer hold in environments where access is exercised continuously by systems, integrations, and automated processes rather than through intermittent human sessions.

Access is now initiated by services, APIs, applications, automated workflows, and machine identities operating across distributed environments. Federated trust relationships, token validation, and service-to-service authentication increasingly underpin critical business operations, extending access beyond traditional user-driven interaction models.

As these interactions span multiple systems and operate continuously beyond direct human oversight, access behaviour becomes more dynamic, variable, and difficult to govern through static identity models alone.[3]

Despite these changes, most identity systems continue to grant access using governance assumptions built around stable roles and predictable access patterns. This creates a structural mismatch between how access is defined and how it accumulates over time.

Modern attacks increasingly exploit trusted identities rather than infrastructure vulnerabilities.

3. Identity Drift and Misalignment

Roles evolve, systems are updated, integrations outlive their original function, and temporary access frequently becomes permanent by default. Over time, governance struggles to keep pace with these changes, resulting in identity drift as permissions accumulate beyond current operational requirements.

Because this misalignment develops gradually, it often remains undetected until a compromise or significant operational change exposes the extent of accumulated access.

When credentials are compromised, attackers inherit this accumulated access, significantly increasing the blast radius of compromise and escalating business impact.[5]

External Document © 2026 Infosys Limited

4. AI as Both Threat and Amplifier

AI and automation introduce a dual reality for identity security.

Defenders use these capabilities to improve anomaly detection, contextual analysis, and response at scale. Attackers use them to automate credential misuse, generate convincing impersonation signals, and rapidly identify overprovisioned access paths.[3]

AI-driven tooling also enables attackers to analyse identity relationships, privilege dependencies, and access pathways at a speed and scale that traditional governance processes were never designed to manage.[6]

These capabilities significantly compress attack timelines and reduce the opportunity for manual intervention. Automated tooling reduces the time required to discover excessive access, exploit accumulated permissions, and move laterally using valid identities.[3]

In environments where access is excessive or poorly aligned, automation converts latent identity risk into immediately exploitable exposure.[3]

5. Real-World Identity Failure Scenarios

These scenarios illustrate a consistent pattern: identity controls are often functioning as designed but are no longer aligned with how access is exercised in modern environments.

Scenario 1: Deepfake-Enabled Identity Reset

An attacker uses AI-generated voice or video to impersonate a user during a service desk interaction. By convincingly replicating identity signals, the attacker successfully bypasses verification processes, resets authentication factors, and gains access to existing permissions.

Risk Exposure

AI-generated impersonation can replicate trusted identity signals such as voice and appearance, making traditional verification methods unreliable. This creates a high-risk entry point that is difficult to detect and increasingly common in targeted attacks.

Control Gap

Identity verification relies on static signals and procedural checks that were not designed to withstand AI-driven impersonation. This highlights the need for stronger identity proofing, contextual validation, and tighter control over high-risk processes such as credential resets and MFA changes.

Scenario 2: Unmanaged Machine Identity Exposure

A service account or API identity is created to support a system integration or migration. The identity is granted elevated permissions to ensure continuity of service but remains active long after its original purpose is complete. An attacker compromises the credentials and uses the identity to access systems without triggering detection.

Risk Exposure

Machine identities often operate continuously, hold persistent access, and lack clear ownership. Because their behaviour aligns with expected system activity, misuse can go undetected for extended periods, creating a persistent and difficult-to-identify attack surface.

Control Gap

Non-human identities are frequently excluded from consistent governance processes. This demonstrates the need for defined ownership, enforced lifecycle controls, restricted permissions, and continuous monitoring to ensure access remains aligned with operational purpose.

Scenario 3: AI-Accelerated Privilege Escalation

An attacker gains initial access through compromised credentials or a synthetic identity. Using automated techniques, they rapidly identify over-provisioned accounts and exploit accumulated permissions to escalate privileges and move laterally across systems.

Risk Exposure

AI enables attackers to analyse access patterns and exploit excessive permissions at speed. Identity drift and over-provisioning significantly increase the impact of compromise, allowing attackers to expand access without resistance.

Control Gap

Access is not continuously aligned to current requirements, allowing permissions to accumulate over time. This highlights the need for ongoing governance, enforcement of least privilege, and continuous correction of access to ensure permissions align with actual usage.

Across each scenario, failure is driven not by infrastructure weakness, but by a control gap and misalignment between identity, access, and behaviour.

External Document © 2026 Infosys Limited

Recommended Solution

Continuous Identity Alignment

In environments where access is exercised continuously, alignment between identity, access, and behaviour does not persist by default.

Continuous Identity Alignment is the operating model required to address this reality. It is the ongoing process of ensuring access remains consistent with organisational intent as identities, systems, and behaviours evolve over time.

Identity must function as an active control system that dynamically evaluates and corrects access based on observed activity and operational context.

Without this level of fidelity, Zero Trust decisions are made on outdated assumptions rather than current use, introducing risk instead of reducing it.[1]

Access doesn't stand still, and identity management can't either.

Core Identity Controls

This operating model is achieved through a set of interconnected control domains that must work together. This continuous alignment occurs throughout the access lifecycle rather than relying solely on verification at login. Behavioural signals, contextual authentication, and machine identity telemetry help reassess trust as identity activity evolves across environments.

Trust establishment addresses whether an identity can still be relied upon when traditional verification signals are no longer sufficient, while access shaping focuses on constraining exposure by aligning permissions to what is required rather than what is convenient.

Lifecycle alignment ensures access changes as roles, systems, and integrations evolve. Behavioural validation introduces context by assessing how access is exercised day-to-day.

From an operational standpoint, these controls cannot be treated independently. Weakness in any one area creates misalignment elsewhere. This is most visible with machine and automated identities, where permissions tend to persist, usage is continuous, and governance is often fragmented.

Without consistent control across these domains, misalignment becomes systemic rather than exceptional.

What Leading Organisations Do Differently

Most organisations still manage identity as a provisioning function. Access is assigned based on a role or request, reviewed periodically, and assumed to remain appropriate between review cycles.

Leading organisations have moved away from assuming access remains appropriate between review cycles after recognising how quickly access relationships drift as operational dependencies evolve.

Rather than treating identity as a static assignment process, these organisations continuously evaluate how access is used, reduce standing privileges, and correct misalignment as it emerges across systems, APIs, and automated workloads.

Periodic governance alone cannot keep pace with the speed at which modern access relationships evolve.

From an IAM perspective, access cannot be treated as a static configuration when it functions dynamically across systems and integrations.

What Good Looks Like

Identity maturity is best understood through an organisation's ability to ensure access continues to reflect real usage over time.

In lower maturity environments, access is defined through static roles, visibility into usage is limited, and permissions accumulate quietly until a compromise reveals how far access has drifted beyond intent.[2]

In more mature environments, access is continuously evaluated against observed behaviour. Least privilege is enforced dynamically, standing access is minimised, and both human and non-human identities are governed through the same control expectations.[1]

The practical result is reduced blast radius, earlier detection of anomalous access, and greater confidence that identity controls reflect how the environment operates.5

External Document © 2026 Infosys Limited

Continuous Identity Control Architecture

Figure 1 illustrates an identity control architecture designed for environments where access is exercised continuously by human, privileged, and machine identities.

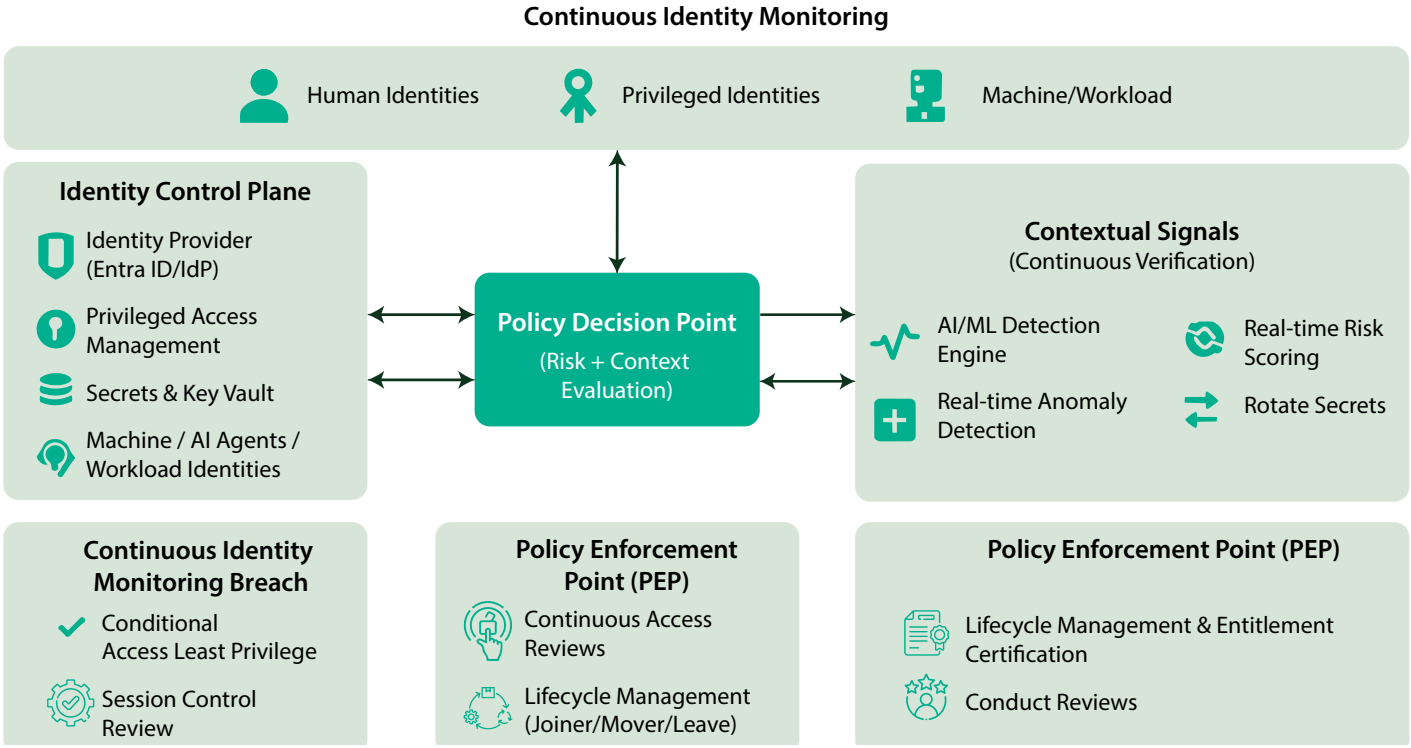


Figure 1. Reference Architecture for Continuous Identity Control

The architecture centres on a Policy Decision Point (PDP), which evaluates identity risk and access context using telemetry from identity systems, behavioural analytics, contextual verification signals, and real-time risk assessment capabilities.

Identity providers, privileged access management platforms, secrets vaults, and machine or workload identity systems form the identity control plane responsible for authentication, privileged access governance, credential management, and machine identity control.

Contextual verification capabilities continuously reassess trust using behavioural analytics, anomaly detection, AI-assisted risk scoring, session evaluation, and secret rotation to identify and respond to suspicious identity activity.

Policy Enforcement Points (PEPs) apply dynamic controls across conditional access, session management, lifecycle governance, entitlement certification, and continuous access review processes to ensure access remains aligned with operational requirements.

This model enables organisations to move beyond static provisioning toward adaptive identity governance aligned with real-world access behaviour and continuously evolving risk conditions.



External Document © 2026 Infosys Limited

Implementing Continuous Identity Control

For IAM teams, these challenges emerge operationally through fragmented visibility, access sprawl, and governance fatigue. Without deliberate control recalibration, governance struggles to keep pace with evolving access dependencies across modern environments.

Transitioning to continuous identity control therefore requires a structured approach that prioritises visibility, access reduction, and ongoing governance. Rather than introducing complexity upfront, leading organisations focus on establishing control across identity environments before layering advanced capabilities.

1. Establish Identity Visibility

Organisations must begin by building a unified view of all identities across the environment.

This includes:

- Workforce identities
- Privileged accounts
- External users
- Service accounts and APIs
- AI-driven and automated identities

Without a unified view of identity activity, it's not possible to determine how access is distributed or where risk exists. In complex environments, a lack of visibility is often the primary barrier to effective identity control.

2. Reduce Standing Privilege

Excessive and persistent access is one of the most common drivers of identity risk.

Elevated access often remains active long after the operational requirement that justified it has changed or expired.

Reducing standing privilege means removing unnecessary access, enforcing least privilege, and limiting elevated permissions to defined time windows. This restricts blast radius during compromise and reduces the ability to move laterally across systems.

3. Strengthen Identity Control Points

Attackers increasingly target operational processes rather than technical vulnerabilities.

High-risk control points include:

- Credential resets
- Access approvals
- Administrative actions

These processes must be secured through stronger verification, contextual validation, and consistent oversight. Weaknesses in these areas are frequently exploited through impersonation and social engineering.

4. Enforce Lifecycle Governance

Access must remain aligned with organisational roles and system requirements over time.

This requires accurate provisioning and deprovisioning, alignment with organisational change and evolving system dependencies, and regular access reviews to ensure permissions reflect current needs. Breakdowns in lifecycle governance are a primary cause of identity drift, where access persists beyond its intended purpose.

5. Introduce Continuous Monitoring

Static controls are insufficient in environments where access activity is continuous and increasingly automated.

Organisations must be able to identify changes in behaviour, privilege use, and access patterns as they occur rather than relying solely on periodic review.

Behavioural analytics and contextual evaluation are required to understand access usage and identify misuse before misalignment becomes material risk.

6. Govern Machine and AI Identities

Non-human identities must be governed with the same level of discipline as user identities.

Unlike workforce identities, machine and automated identities often operate continuously, authenticate programmatically, and persist across multiple systems without direct human oversight.

This includes:

- Assigning clear ownership
- Limiting permissions
- Enforcing lifecycle controls
- Monitoring usage and behaviour

Across modern enterprise environments, machine and AI-driven identities represent the least visible yet highest-risk attack surface due to their persistence and elevated access.

Effective identity security depends on maintaining alignment between identity governance, access activity, and observed behaviour.

Organisations that establish visibility, actively manage access, and enforce governance consistently are better positioned to control identity risk in AI-driven environments.

External Document © 2026 Infosys Limited

Conclusion and The Future of Identity Security

Identity systems define how access is established, exercised, and retained across modern enterprise environments.

Attackers increasingly exploit these systems directly, using credential compromise, impersonation, and misuse of legitimate access rather than infrastructure vulnerabilities. Because these actions align with valid access pathways, they are harder to distinguish from normal activity and more difficult to contain.[2]

As access increasingly operates through systems, integrations, APIs, and automated processes, the limitations of identity models built for predictable, human-driven behaviour become increasingly visible. In these conditions, fixed access definitions and periodic review are no longer sufficient to maintain control.[3]

Maintaining effective identity security now depends on maintaining alignment between identity, access, and observed behaviour. This requires organisations to move beyond governance models that assume access remains appropriate after provisioning and instead operate identity as an adaptive control system capable of responding to evolving operational context and risk.

As AI accelerates the scale and speed of identity-driven attacks, organisations that continuously align identity governance, access relationships, and operational behaviour will be better positioned to contain compromise, reduce blast radius, and maintain trust across distributed environments.



External Document © 2026 Infosys Limited

References

1. NIST SP 800207: Zero Trust Architecture, 2020, https://csrc.nist.gov/publications/detail/sp/800-207/final
2. Verizon Data Breach Investigations Report, 2025, https://www.verizon.com/business/resources/reports/dbir/
3. Microsoft Digital Defense Report, 2025, https://www.microsoft.com/security/security-insider/threat-landscape
4. ENISA Threat Landscape, 2024, https://www.enisa.europa.eu/topics/threat-risk-management/threats-and-trends
5. IBM Cost of a Data Breach Report, 2024, https://www.ibm.com/reports/data-breach
6. NIST AI Risk Management Framework, 2023, https://www.nist.gov/itl/ai-risk-management-framework

Additional Industry References

- OWASP API Security Project, 2024, https://owasp.org/www-project-api-security/
- OWASP Top 10 for Large Language Model Applications, 2025, https://owasp.org/www-project-top-10-for-large-language-model-applications/
- World Economic Forum – Digital Identity and Trust, 2024, https://www.weforum.org/topics/digital-identity/
- Sysdig Cloud Native Security & Usage Report, 2025, https://www.sysdig.com/resource/cloud-native-security-and-usage-report

About the Author



Ruchit Deshpande

Security Solutions Director, The Missing Link

Ruchit leads the Security Architecture, Engineering and Project Management teams, with 15+ years of experience in Cyber Security. He has technically led successful deliveries of program of works for cyber posture uplifts, assessments and implementation services for Enterprises across multiple industries. Ruchit is adept at presenting/speaking at Board level and C-suite audiences.

For more information, contact askus@infosys.com

Infosys
Navigate your next

© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.