The Identities Nobody Put on the Org Chart

Why Non-Human Identity Is an Operating Discipline, Not Another Security Tool

AHEAD

Software now calls APIs, deploys infrastructure, moves data, renews certificates, and operates production systems without a person at the keyboard. Each action depends on identity.

That identity may be a service account, cloud role, API token, certificate, CI/CD credential, automation bot, or AI agent. These identities are how modern environments operate, but most identity programs were designed around people.

A sound non-human identity (NHI) strategy is not a larger vault or dashboard. It is an operating model that can answer five questions for every machine actor: What is it? What can it do? Who is accountable for the service? How is access issued and removed? How quickly can it be stopped when its behavior no longer matches its purpose?

Non-Human Identity, Defined

A non-human identity is an identity used by software, infrastructure, a device, or an automated process rather than by a person. It is what an automated actor presents before it is allowed to act.

Infrastructure teams may say *machine identity*. Cloud teams may say *workload identity.* Directory teams may say *service account*. Developers may say *API key, token, or service principal.* Public-sector guidance often uses *non-person entity*.

No matter the term, the structure is the same:

- The actor: which application, workload, process, bot, device, pipeline, or agent is acting?
- The authenticator: what proves its identity — a password, token, key, certificate, or signed assertion?
- The entitlement: what can it reach, read, change, deploy, sign, or trigger?

A practical program covers service accounts, workloads, bots, automation, API keys, tokens, certificates, cryptographic keys, and AI agents. The authenticator is not the identity; risk comes from the relationship between the actor, the authenticator, and the access it receives. NHI is broader than AI, although AI makes the issue more visible by combining delegated authority, tool access, and autonomous action.

Why The Issue Is Urgent

The scale is changing faster than the operating model. CyberArk's 2025 research reported 82 machine identities for every human identity, while a separate study found that 79% of organizations expect machine identities to grow over the next year — some by as much as 150%. Half reported an incident or breach linked to a compromised machine identity, and 42% lacked a cohesive strategy. These are survey findings, but they show identity populations expanding while ownership and controls remain fragmented. See the 2025 Identity Security Landscape and 2025 State of Machine Identity Security Report.

NIST's Zero Trust Architecture explicitly includes people and non-person entities, credentials, access management, operations, and hosting environments. Its requirements are direct: verify explicitly, restrict access to what the subject needs, and grant minimum necessary privilege. CISA applies the same logic to entities through just-in-time access, orchestration, and automated identity analytics.

NHI is a shared responsibility across identity, security operations, engineering, cloud, PKI, DevOps, audit, and the business. The consequences fall into three connected areas: exposure to external threats, friction in business and AI adoption, and operational or cyber-resilience failure.

The Attack Paths Are Also Business Paths

Enterprise AI security challenges are rooted in the growing complexity of modern data estates. Sensitive information is distributed across cloud platforms, SaaS applications, collaboration tools, databases, file repositories, archives, and legacy systems, often without consistent visibility, ownership, or classification. As AI becomes embedded within business processes, these longstanding data management challenges become more consequential because AI operates on the information already available within the environment. Understanding the relationship between data sprawl, classification, governance, and access is therefore critical to managing AI-related risk.

1. External threats use valid machine access

Attackers do not need to compromise an AI model or defeat a person's MFA if they obtain a valid machine credential. A token in code, a privileged service account, a misbound cloud role, or a signing key in a build system can provide access that looks legitimate downstream.

Common exposures include long-lived accounts, copied trust policies, CI/CD identities with production access, and certificates or keys that are stolen, misissued, renewed after retirement, or allowed to expire without dependency knowledge. AI agents add delegated user context, tool permissions, and downstream credentials at machine speed.

NIST SP 800-207A says workload identity should be continuously verifiable. CISA's 2025 warning about the compromised *tj-actions/changed-files* GitHub Action shows why pipeline identity, dependency trust, secret scope, and revocation must be managed together.

The response is to remove durable credentials where possible, reduce standing privilege, narrow trust, monitor behavior, and test revocation.

2. Weak NHI capability can slow growth and halt AI adoption

Cloud platforms, partner ecosystems, digital products, and AI workflows depend on reliable application-to-application access. Manual identity issuance creates queues; broad access creates risk that eventually blocks delivery.

AI efforts often stall when an experiment needs production access because teams cannot answer who sponsors the agent, what data and tools it can reach, whether credentials are scoped, what it did, or how to stop it without affecting unrelated services. A sound NHI capability turns that debate into a design question: what identity, authorization, telemetry, and cutoff pattern does the use case require?

The same capability supports acquisitions, customer-facing services, and self-service platforms by making access repeatable, scoped, observable, and reversible.

3. NHI is an operational and cyber resilience dependency

NHI failures can produce outages, failed deployments, broken integrations, or an inability to restore systems after an attack. A credential change can break an application when consumers are unknown; a certificate can expire when dependencies are unmapped; a revoked token can stop a release process when its replacement path is untested.

Resilience requires an authoritative inventory, dependency mapping, versioned identity configuration, tested rotation and revocation, and a documented order for restoring trust. Recovery teams should know which identities are required to re-establish the identity plane, which are safe to re-enable, and which must remain disabled.

Exercise the controls: test revocation and recovery time, validate certificate replacement, rebuild policies from known-good state, send machine activity to security operations, and include service accounts, workloads, certificates, keys, and agents in recovery runbooks.

Mature NHI programs help the business deploy safely, change systems without guessing, detect misuse, and recover without reintroducing compromised identities.

Why Assigning An Owner Is Not Enough

Most NHI guidance says every identity needs a named owner. That is necessary, but also incomplete. A person attached to a record does not prove they understand the identity, can change its permissions, or can safely shut it down. It also fails for workloads created and destroyed with deployments.

A workable model separates accountability:

- Executive sponsor: priority, risk tolerance, and funding.
- Program owner: standards, metrics, exceptions, and cross-team decision rights.
- Application or service owner: purpose, impact, dependencies, and retirement.
- Technical custodian: creation, attachment, rotation, and removal in the platform.
- Security operations and risk: monitoring, containment, and evidence.

For ephemeral workloads, accountability should come from deployment context — application catalog, repository, namespace, environment, cost center, on-call team, and business service. The workload inherits accountability from the service definition that created it, and the technical custodian operates the control, but does not replace business accountability.

This matters most during reorganizations. When identities are “owned” by former employees, dissolved teams, or executives who never operated the service, shared accountability allows you to re-establish the service relationship, identify a technical custodian, confirm purpose and dependencies, and either remediate or retire the identity more immediately.

The crux of the question must shift from who the owner is to *who is accountable, who can operate the control, and what evidence proves the access is still needed?*

A Practical Strategy Beyond Tools

Establish one control model

Human, machine, and AI identities need different implementations but the same principles: know it, own it, limit it, watch it, stop it, and measure it. Access decisions should go beyond credential validity and include workload, environment, resource, behavior, and business purpose.

Build inventory and prioritize consequence

Join evidence from directories, cloud IAM, Kubernetes, CI/CD, repositories, secrets managers, PAM, PKI, CMDB, application catalogs, ITSM, and security logs. Identify the consumer, runtime, entitlement, credential, accountable service, and disablement impact.

Start with production control planes, privileged cloud roles, build systems, signing infrastructure, sensitive data paths, and externally exposed services. Define normal behavior, minimum permissions, and tested revocation time for each path.

Use the safest mechanism and automate it

Remove durable credentials with managed identity or federation. Shorten credential life with short-lived certificates, task-scoped tokens, or just-in-time access. Vault and rotate static credentials that cannot yet be removed. And contain and monitor legacy exceptions with a retirement date.

NHI controls belong at creation, deployment, permission change, ownership change, retirement, and incident response. Policy-as-code and infrastructure-as-code can version definitions, test changes, detect drift, and rebuild known-good state. AI can correlate records and rank risk, but humans should retain policy, exceptions, high-impact approvals, and the kill decision.

To know that the strategy is working, we can measure inventory coverage, accountable owners for high-risk identities, privileged paths with tested revocation, managed or short-lived credential use, credential age, and time to contain anomalies.

A 90-Day Starting Point

First 30–60 days

- Name an executive sponsor and NHI product owner.
- Define in-scope populations and the minimum record for each.
- Choose one high-risk machine path and map every credential consumer.
- List AI agents with sponsor, purpose, and scope.
- Test disablement, revocation, rollback, and service recovery.

Next one to two quarters

- Remediate by application and service wave.
- Remove durable credentials where federation or managed identity is available.
- Move remaining high-risk secrets into managed rotation and just-in-time access.
- Connect identity telemetry to security operations and ITSM evidence.
- Automate certificate renewal and retirement based on service context.

Tool selection belongs inside this sequence. Keep and refit platforms that can support the model, and add capabilities when the gap is discovery, workload identity, certificate automation, or agent governance.

AHEAD's NHI Practice

AHEAD treats NHI strategy, execution, and optimization as a core identity competency — not an attachment to workforce IAM and not a reason to sell a single product.

The AHEAD Identity team includes multiple leaders with more than 15 years of average experience across the principal layer. That depth spans directory and governance, PAM, PKI, cloud infrastructure, platform engineering, application delivery, security operations, and recovery. AHEAD's growing consulting practice works across those disciplines because NHI problems rarely stay within one tool or team.

Our work involves three connected motions:

- Advise: establish the inventory, risk model, decision rights, target architecture, and delivery roadmap.
- Build: connect identities to applications and workloads, remediate high-consequence paths, and integrate the platforms already in use.
- Operate and optimize: measure coverage and revocation, automate lifecycle events, run recovery exercises, and improve the model as infrastructure and AI adoption change.

AHEAD maintains strong partnerships across privileged access, secrets, certificate lifecycle, cloud and workload identity, and NHI and AI-agent governance. Technology choices follow the control model and evidence, not the other way around.

The delivery model is designed for machine scale. Identity policies, roles, integrations, and control definitions can be versioned as code that is validated in pipelines, checked for drift, and connected through infrastructure and orchestration. Governed AI-assisted analysis can correlate data and rank risk. The outcome is a repeatable identity capability that helps clients reduce exposure while continuing to ship, automate, and recover.

Final Thoughts

Non-human identity security is not a future problem created by AI. It is the accumulated identity population of modern infrastructure, made more urgent by cloud scale, shorter lifecycles, automation, and agentic software.

The practical response is to establish one control model, derive accountability from service context, remove durable secrets where possible, enforce policy where identities are created, connect telemetry to response, and measure whether risk and recovery time improve.

Mature programs make machine identity explainable, bounded, observable, and stoppable. That is the standard an identity program should meet before it gives software authority to act on the business's behalf.

AHEAD

Combining cloud-native capabilities in software and data engineering with an unparalleled track record of modernizing infrastructure, we're uniquely positioned to help accelerate the promise of digital transformation.

Visit us at [ahead.com.](ahead.com)

Global Hub Offices

CHICAGO

444 W. Lake Street
Suite 3000
Chicago, IL 60606
United States

NEW YORK

500 5th Avenue
17th Floor
New York, NY 10010
United States

UNITED KINGDOM

310 Wharfedale Rd
Winnersh, Wokingham
RG41 5TP
United Kingdom

NETHERLANDS

Kerkenbos 10-123
6546 BJ Nijmegen
The Netherlands

INDIA

L-21, One Horizon Center
Golf Course Road, DLF Phase V
Sector 43, Gurugram 122002
India