

WHITE PAPER

Securing non-human identities

A complete strategy for identity resilience across hybrid and AI-driven environments

Introduction

Non-human identities now outnumber human identities by as much as 82:1. And with cloud automation and agentic AI, that number will only continue to grow. These identities quietly power modern environments, operating continuously in the background, authenticating silently, and often carrying far more privilege than their narrow function requires.

Quest Software's latest identity threat detection and response (ITDR) research found that non-human identities are the most difficult area of identity infrastructure to secure, ranking ahead of privileged accounts, critical Tier 0 assets, cloud identities, and on-premises AD and legacy systems. That finding is especially significant given how deeply embedded non-human identities are across all these layers.

What are non-human identities (NHIs)? What impact do they have on cybersecurity? Why are organizations struggling with them? And what does a modern non-human identity security strategy look like?

Non-human identities: what they are and why they matter for security

Defining non-human identities

Non-human identities, also referred to as workload identities or programmatic entities, are identities not attached to a human user. They exist to enable systems, applications, services, and devices to authenticate and operate autonomously.

Under Gartner's NHI and machine IAM taxonomy, non-human identities span partners, customers, and machines. Within the machine category, NHIs include:

- Devices – Desktop, mobile, and IoT
- Workloads – VMs, containers, service accounts, and applications (including AI agents)

As a distinct identity class, non-human identities behave differently from human users. They:

- Perform administrative and scheduled tasks autonomously
- Operate programmatically, running scripts or programs continuously
- Authenticate using secrets like API keys, tokens, and digital certificates
- Perform a singular, highly specific function
- Spin up and tear down quickly within DevOps pipelines, making manual tracking difficult

	Human identities	Non-human identities
Primary owner	A person – employee, contractor, consultant, etc.	A resource, workload, bot, service, or AI agent
Authentication	Interactive – passwords, MFA, biometrics	Static and silent – API keys, tokens, certificates, shared secrets
Lifecycle	Joiner-mover-leaver process, managed by HR	Ephemeral or infinite; spun up instantly for a task or left for years
Behavior	Interactive, variable, usually operate during business hours	High-volume, repetitive, predictable patterns, operate 24/7
Risk profile	Phishing, weak passwords, social engineering	Overprovisioning, hard-coded secrets, secret sprawl
Volume	Linear growth (one per person)	Exponential growth (dozens per app/microservice)

Non-human identities vs. human identities

Non-human identities work differently from human identities in almost every way.

Human identities are interactive and authenticate with controls like passwords and MFA. They are governed through established joiner-mover-leaver processes, and their behavior varies, but is usually tied to business hours.

Non-human identities operate very differently. They authenticate silently using static credentials like API keys, tokens, certificates, or shared secrets. They often persist indefinitely without formal lifecycle management and run continuously and programmatically, executing repetitive tasks at high volume.

These differences fundamentally change their risk. A compromised human identity triggers alerts due to unusual login behavior or unexpected access patterns. A compromised non-human identity does not. Its activity appears legitimate because it is legitimate, using the same credentials and executing the same actions, only under an attacker's control.

Non-human identities grow exponentially, with dozens created for every application or service. Without governance, they accumulate quickly and remain largely invisible.

AI agents: the next evolution of non-human identities

The NHI population is already large and poorly governed. AI agents represent their next evolution. Like service accounts, AI agents act as themselves, performing autonomous functions in the background, like scheduled tasks, maintenance, or administration. However, unlike service accounts running a defined script on a schedule, AI agents can make decisions and act independently across environments.

Governing these identities demands continuous visibility into what an agent is doing, why the agent is doing it, and whether the agent's actions align with its intended purpose. Privileges must be tightly constrained and adjusted as behavior evolves.

The challenges of securing non-human identities

Lifecycle management

Human identities are tracked and managed by HR, but non-human identities have no resources to manage their lifecycle. They are created by developers or automated scripts for a specific task, and many are overprivileged for convenience during testing. These NHIs are designed to be temporary, but when cleanup never happens, they become what security teams call a "zombie identity," sticking around in the environment indefinitely.

The lifecycle management problem is further compounded by secret sprawl. More non-human identities than human identities mean more secrets to manage, especially when maintaining good hygiene and rotating credentials.

Authentication

There is no equivalent of multi-factor authentication for non-human identities. Authentication is static and silent, relying on API keys, tokens, certificates, and shared secrets. Once compromised, these credentials provide unrestricted, programmatic access without triggering security alerts and can operate inside environments as a trusted identity.

Overprivilege and hard-coded secrets

Overprovisioning is endemic in non-human identity environments. Developers frequently assign broad permissions during testing to ensure functionality, then never trim those permissions back. The result is NHIs with access far exceeding what their function requires, which dramatically increases blast radius when compromise occurs.

Hard-coded secrets compound the problem. Private keys compiled into device firmware, credentials stored in plaintext in source code, and certificates with multi-year validity that have never been reviewed all represent exposure that accumulates over time without governance.

Legacy systems

Many organizations still rely on legacy platforms such as mainframes and AS/400 systems that lack automatic rotation. When a critical business process depends on an integration that has run without issue for years, there is little incentive to modernize. Certificates can remain valid for years, carrying overly broad privileges far longer than necessary and eventually expiring without warning, causing outages.

Visibility

In large organizations, different business units operate in silos. Each has developers using AI and NHIs to build applications. From a central IT perspective, that equates to potentially thousands of NHIs in the environment with different permissions, siloed across the business. Organizations are struggling to gain visibility and governance over them; there are simply not enough human hours to manually review and monitor thousands of automated identities.

Ownership

When an incident occurs, organizations need to be able to quickly determine who is responsible for a given identity, what function it serves, and whether its behavior is expected or anomalous. Without a defined owner, environments lose valuable time as security teams sift through logs and activities to understand the purpose and scope of a compromised identity.

The absence of ownership also undermines effective response and accountability. Non-human identities often persist across teams and business units, outliving the processes they were created to support. When responsibility is unclear, remediation slows and risk increases over time. For regulated and publicly traded organizations, the inability to explain how an identity was governed can have implications that extend to executive leadership.

Third-party risk

When access is granted to third-party services, NHI credentials become the mechanism of trust between environments. A compromise at the vendor level can therefore propagate to the customer without interacting with a human user. This interconnected use of NHIs creates supply chain risk that is difficult to inventory and even harder to govern. Over time, these trusted connections accumulate and expand the attack surface beyond control.

How NHIs are being exploited

Midnight Blizzard

In the Midnight Blizzard nation-state attack, state-sponsored actors gained access through a legacy application in a non-production environment that had been granted elevated permissions within a Microsoft tenant. The application had been created for a legitimate purpose, but over time, it became dormant. It was not actively used or monitored. Because the account retained its permissions and authenticated silently, attackers could use it to access executive emails undetected.

Storm-0501

Storm-0501 began with ransomware campaigns across U.S. school districts in 2021. The cybercriminal group evolved to target healthcare data in 2023, and by 2024 had pivoted to using NHIs to target cloud environments.

Most organizations rely on synchronization processes that connect on-premises Active Directory to Entra ID using a service account with elevated permissions. Storm-0501 compromised one of these accounts, which had not been updated with current security best practices, then elevated the privileges of a standard user account, created a new directory, and extracted the Active Directory database file.

The initial compromise did not involve a human user account, nor did it require bypassing multi-factor authentication. Instead, it exploited the inherent trust placed in a non-human identity running a critical background process.

Snowflake data breach

In the Snowflake campaign, attackers leveraged infostealer malware to harvest credentials for service accounts that lacked MFA. Once obtained, these static credentials provided silent, programmatic access that appeared legitimate. Instead of reaching customers directly, attackers exploited trusted service relationships, using compromised credentials to exfiltrate data from massive customers like Ticketmaster and Santander.

Dropbox Sign incident

In the Dropbox Sign breach, attackers compromised a backend automated service account that had continuous access to customer data. Because the account's access was static and silent, there was no need for attackers to create a new account or change any permissions. Once the credential was obtained, access appeared normal. Detection typically depends on changes in behavior, and because behavior appeared normal, the compromise went undetected far longer than a traditional account takeover would have.

What these attacks have in common

While the specific details of these incidents differ, the underlying conditions do not. In every case:

- A non-human identity was created for a legitimate purpose
- Permissions were broader than its function required
- Credentials persisted without rotation or review
- Monitoring and ownership were minimal or absent

These are not sophisticated failures. They are gaps in basic hygiene, and they are common. As organizations expand their use of automation and cloud services, the number of non-human identities grows. Without deliberate visibility, governance, and lifecycle management, each new NHI is a low-friction entry point for attackers seeking persistent, stealthy access.

Best practices for managing and securing non-human identities

1. Discovery and inventory

Protection depends on visibility. Organizations need automated solutions to detect and identify NHIs. At a ratio of 82:1, manual review is not feasible. Discovery tooling must be automatic, programmatic, and ideally AI-assisted to match the scale and speed of modern adversaries.

Categorization is equally important. In regulated industries, certain accounts need to be handled differently because of requirements or regulations, and those accounts need to be identified. With the right classification, it's very easy to provide information to an auditor and identify what permissions exist and what they're being used for.

2. Lifecycle automation

Typically, non-human identities have no lifecycle management and no rotation. Those are standing secrets that last forever. If they're compromised, they can provide persistent access and do significant damage. Implementing secret rotation with keys rotated every 30-90 days addresses the majority of NHI lifecycle issues.

3. Least privilege monitoring

A least privilege model applies even more to non-human identities than to human ones, because NHIs have a narrow, defined set of tasks to execute. Their behavior is programmatic and predictable. Even for agentic AI, there is typically a very specific set of search or execution tasks. Privileges can be tightened precisely, especially with a tool that analyzes access and surfaces overprovisioned identities.

Visibility into what an NHI is doing, what pattern it follows, and what is changing enables a virtuous cycle. Security teams can implement the narrowest possible privileges, review what the identity is doing against what it is permitted to do, and tighten access further. Permissions that are assigned but never used represent unnecessary exposure and should be removed.

4. Identity threat detection and response

In the age of agentic AI, reactive response is too slow. Programmatic threats execute faster than any human-led response, and damage can be done in seconds. Organizations must adopt proactive, preemptive protection for their most sensitive assets and crown jewel objects.

NHIs operate continuously around the clock, without the business hours that typically anchor anomaly detection. Constant monitoring for anomalous activity is essential, as is AI-assisted investigation and remediation. Because when compromise causes an impact within the directory, organizations need automated, safe, and fast recovery.

How Quest can help

Quest Identity Defense

Quest Identity Defense is an AI-powered, hybrid Active Directory and Entra ID security solution that continuously identifies identity risk, blocks unauthorized changes to critical Tier 0 assets, and instantly contains attacks before they spread. It is a core component of the Quest Security Management Platform, which unifies identity threat detection, recovery, and secure migration into a single platform.

The Quest Security Management Platform aligns to all five pillars of the NIST Cybersecurity Framework: identify, protect, detect, respond, and recover. As ITDR has evolved beyond a detection-only focus, organizations need to coordinate security, resilience, and recovery without juggling multiple point solutions to protect their most important assets.

Quest Identity Defense delivers:

Entra ID workload analysis

- Risk analysis – Scores each application or service principal based on the reach of its permissions
- Ownership identification – Surfaces whether an application is owned by a group or individual user
- Certificate and secret monitoring – Monitors certificates and secrets to show how each application is protected and when that protection expires
- Permission visibility – Identifies if an API has delegated permissions (signing on behalf of a user) or application permissions (acting with its own direct access)

NHI identification for Active Directory

Active Directory environments that have been in place for 25 years or more may contain service accounts, managed service accounts (MSAs), group managed service accounts (gMSAs), delegated managed service accounts (dMSAs), and computer accounts that have accumulated over decades, many of which have never been reviewed or classified.

Quest Identity Defense applies a customizable AI engine to this challenge, using configurable rules like naming conventions and SPN presence to identify and classify NHIs within Active Directory. Each identity receives a confidence score indicating the likelihood that it is an NHI, with the ability for administrators to manually verify, recertify, or add identities. When accounts are categorized in accordance with HIPAA, PCI, or SOX requirements, organizations can respond to audits through reporting rather than last-minute discovery.

Proactive protection of crown jewel assets

Quest Identity Defense goes beyond detection and alerting to preemptive protection. Dynamic Shields Up containment stops attacks mid-flight and freezes critical objects to prevent compromise, limiting attacker lateral movement. Modification of the most sensitive identities and objects is restricted to only the identities explicitly trusted to make those changes. Even if a credential is obtained, the ability to use it to modify critical assets is constrained.

Deep visibility, auditing, and intelligent detection

Quest Identity Defense provides deep visibility and auditing for AD and Entra ID, with high-quality signals and intelligent context that reduces noise and enables faster investigation. For NHIs, this means monitoring the behavior of each identity, flagging anomalous activity that deviates from expected patterns, and providing AI-assisted investigation capabilities that accelerate response.

> "Quest Identity Defense is the best tool we could find available for identity threat hunting."
>
> CISO, Large Media Company

Recovery and resilience

A mature identity security program must account for what happens when prevention and detection controls are not enough. The Quest Security Management Platform pairs Quest Identity Defense with Quest Identity Recovery, which restores hybrid identity environments up to 90 percent faster after ransomware attacks and includes always-ready Standby Active Directory Forest provisioning. Organizations that address the full identity security lifecycle reduce breach likelihood and are materially better positioned to restore operations quickly when incidents do occur.

About Quest Software

Quest Software creates technology and solutions that build the foundation for enterprise AI. Focused on data management and governance, cybersecurity and platform modernization, Quest helps organizations address their most pressing challenges and make the promise of AI a reality. Around the globe, more than 45,000 companies including over 90% of the Fortune 500 count on Quest Software. For more information, visit www.quest.com or follow Quest Software on X (formerly Twitter) and LinkedIn.

Explore our solutions →

© 2026 Quest Software Inc. I ALL RIGHTS RESERVED.
Quest, and the Quest logo are trademarks of Quest Software Inc.
For a complete list of Quest marks, visit https://www.quest.com/legal/trademark-information.aspx. All other trademarks are properties of their respective owners. Whitepaper-Securing-Non-Human-ID-HO 102485