CYBER THREAT INTELLIGENCE

CRITICAL**START®** Threat Research

TLP CLEAR // [CS-TR-26-0401] The Hidden Threat of Unmanaged Machine Identities in Enterprises

Executive Summary

Enterprises that cannot inventory, constrain, monitor, and rotate its non-human identities (NHIs) are delegating trust to insiders they do not govern. Non-human identities (NHIs), also referred to as machine identities are a dominant population in enterprise environments and are crucial to the operations of small and large-scale organizations. According to Obsidian Security's 2026 NHI guide, NHIs outnumber human identities by 25 to 50x in modern enterprises. NHIs appear as keys, tokens, certificates, secrets, CI/CD and other ephemeral workload credentials[2], and recently AI agents which enable human-system and system-system interactions. Machine identities programmatically authenticate to cloud control planes, source code repositories, SaaS applications, internal API endpoints, data stores, message queues, and infrastructure layers that no human touches directly. That operational invisibility has made them central to modern attack chains.

Non-human identities are exploited for cyberattacks that result in data breaches. Per Obsidian Security, 68% of IT security incidents now involve machine identities[1]. Incidents such as the 2023 Cloudflare breach[12], and attacks on SoundCloud and Betterment through compromised Okta single sign-on environments by Lapsus$ Group[3,5,6] demonstrate how identity access can be leveraged to move across systems and exfiltrate sensitive data. To mitigate broader identity risks, enterprises developed identity governance programs around users. Attackers adapted by targeting the identities that never log off or go on vacations. This report examines the full scope of the NHI attack surface, documents the breach pattern that follows from it, and provides mitigation strategies that address the governance gap at its root.

Introduction

Non-human identities (NHIs) are central to automation and modern agent-driven workflows in enterprise environments. According to CyberArk, NHIs are digital entities used to identify, authenticate, and authorize machines, devices, IT infrastructure, applications, cloud workloads, and automated processes[2]. In practical terms, they include any identity that is not associated with a human user[1]. This category is broader than commonly assumed and limiting it to service accounts overlooks a significant portion of the enterprise attack surface.

Service accounts are often the starting point for understanding NHIs. These operating system and application-level credentials run background services, connect web servers to databases, and execute scheduled tasks. They represent the earliest and most familiar form of non-human identity and continue to play a prominent role in enterprise environments. As organizations adopted SaaS platforms and cloud integrations, NHIs expanded to include API tokens and OAuth applications. These identities enable machine-to-machine communication across systems and services. Unlike human users, they are rarely protected by multi-factor authentication, which allows compromised credentials to provide persistent and direct access.

In cloud-native environments, NHIs have become increasingly dynamic and short-lived. Workload and cloud identities, such as IAM roles assigned to containers, microservices, and serverless functions like AWS Lambda and Azure Functions, are created programmatically to grant services access to resources. Service accounts are typically long-lived and manually managed, while these newer identities can be generated and retired within seconds, often without direct human interaction. The function remains consistent, which is to grant access to systems, but the scale and speed of their lifecycle introduce new complexity for governance[1].

Machine identities form another critical category. These include TLS and SSL certificates, code-signing keys, and device credentials that establish trust between systems rather than representing workloads themselves. They provide the cryptographic foundation that enables secure communication across all other identity types. Operational tooling adds an additional dimension. CI/CD build agents such as GitLab Runners and GitHub Actions operate on behalf of developers and frequently have permission to deploy code into production environments. While they are a form of workload identity, their position in the software supply chain and their level of access make them particularly sensitive.

Other forms of NHIs extend beyond traditional IT systems. Robotic process automation bots replicate human actions across legacy environments and often operate with broad access and limited visibility. IoT device credentials authenticate physical and virtual devices across industries such as manufacturing and healthcare. Industry analysis from Palo Alto Networks and CrowdStrike indicates that each of these identity types introduces distinct security considerations that do not align directly with traditional human identity governance models.

The most recent evolution of NHIs is the emergence of autonomous AI agents[1]. These systems build on existing identity types by using API tokens, assuming workload identities, interacting with services, and in some cases initiating CI/CD processes. The defining characteristic is their level of autonomy. These agents can operate independently and at machine speed across the systems their credentials can access. Autonomous AI agents are explored further in part 2 of this series.

This progression reflects a clear shift in how identities function within enterprise environments. NHIs have evolved from static, long-lived credentials into dynamic and interconnected entities that operate continuously across systems. The focus is no longer on any single identity type, but on the collective impact of identities that enable broad access, operate at scale, and support increasingly autonomous workflows.

The Expanding NHI Attack Surface Built on Trust

As enterprises increasingly rely on non-human identities (NHIs) for automation, cloud workloads, SaaS integrations, and backend services, the number of these identities grows significantly. While these identities are essential to system operations, they are often created and managed by engineering teams with limited visibility from security and governance functions. This gap results in rapid, unmanaged growth that expands the enterprise attack surface.

Recent enterprise incidents highlight how NHIs can play a central role in security events when governance does not fully keep pace with their growth. Attackers increasingly recognize that service accounts and API keys offer durable and programmatic access to systems. Campaigns such as GlassWorm demonstrate this shift[7]. In this case, stolen GitHub tokens were used to force-push malicious code into Python repositories, bypassing human-centric review workflows. The incident illustrates how write access granted to a non-human identity can directly influence critical systems when appropriate controls are not consistently applied.

The software supply chain reflects a similar pattern. Malicious npm packages delivering PylangGhost RAT demonstrate how trust placed in repositories, package managers, and automation pipelines can be exploited through non-human authentication paths[8]. CI/CD systems, dependency publishing workflows, package signing mechanisms, and repository automation all depend on NHIs. A compromise at this layer turns a machine identity into an attacker-controlled insider with direct access to production systems and downstream environments. This illustrates that the non-human identity problem cannot be reduced to secrets management alone. Secrets are simply the artifact. The real issue lies in the governance model behind them: Who owns the identity? What access does it have? How long does it persist? Are its permissions justified? Can anyone distinguish normal operations from abuse? These are questions most organizations address for human employees, but far fewer answer them for the identities running the enterprise.

Further, the proliferation of unmanaged NHIs increases the likelihood of exposing enterprise secrets. GitGuardian reported an 81% increase in AI-service leaks, identifying 29 million exposed secrets on public GitHub[6,15]. The key issue here is not merely the leakage of secrets, but what those secrets represent. These secrets are not just configuration errors; they are enterprise identities with persistent access to critical infrastructure, APIs, and downstream systems. When attackers acquire these credentials, they do not need to compromise human users. Instead, they inherit trust already embedded in the system, gaining privileged access without triggering typical security alerts.

Privilege concentration further amplifies the risk. Research cited by CSO Online, based on Entro Security findings, shows that 97% of NHIs have excessive permissions[14]. Access is also highly concentrated, with 0.01% of machine identities controlling 80% of cloud resources. Compromising one of these identities provides immediate and extensive access without requiring lateral movement. This exposure is compounded by weak credential lifecycle practices. With 71% of NHIs not rotated within recommended timeframes, credentials remain valid far longer than necessary, increasing the window for exploitation.

These conditions define an attack surface built on implicit trust, persistent access, and limited oversight. NHIs operate at scale across critical systems, yet they are not governed with the same consistency or rigor as human identities. The result is an environment where access is widely distributed, difficult to attribute, and readily exploitable. To understand how these risks materialize in practice, the next section examines key breaches where NHIs played a central role.

Unmanaged Machine Identities Enable Enterprise Compromise

The breach reporting conventions of the security industry have a consistent blind spot. Incidents are usually named after the organization affected, such as the Okta breach, the Cloudflare breach, or the Treasury breach, rather than the root cause that enabled them. The result is a fragmented narrative of isolated incidents when the underlying pattern is unified and repeating. Across major enterprise compromises of the past several years, a single structural failure appears in every initial access chain: a machine credential that was over-privileged, unmonitored, or not rotated.

The NHI Management Group's compilation of 40 NHI breaches documents this pattern in aggregate. The incidents below illustrate it in detail[9].

In October 2023, Okta's support case management system was breached through a compromised service account with production support access and no behavioral monitoring. This account provided the defining entry point for what became one of the most consequential supply chain attacks of the year. Files uploaded by 134 Okta customers were accessed, including those of BeyondTrust, Cloudflare, and 1Password[9]. The root cause was a single non-human identity with broad permissions and no governance controls on its use.

The downstream effect at Cloudflare in November 2023 illustrates what poor NHI governance costs in practice[13]. Cloudflare's security team responded to the Okta exposure by rotating thousands of production credentials, but they missed four that were assessed as inactive because there were no recent usage records. These included one token for Moveworks and three service account credentials for Smartsheet, Atlassian Jira, and a third integration. Although labeled inactive, these credentials still had permissions. Using them, attackers accessed Cloudflare's full Atlassian environment, including more than 14,000 wiki pages, 2 million Jira tickets, and nearly 12,000 Bitbucket repositories[16]. Threat actors exploited harvested internal information to attempt pivots into cloud and on-premises environments.

Other incidents follow the same pattern. In January 2024, the Russia-linked APT29 group leveraged a legacy test tenant account at Microsoft that had no MFA, was unmonitored, and remained active long after its intended use[9]. In April 2024, Sisense reported attackers exfiltrating tokens, API keys, and certificates from self-managed repositories lacking production-level controls. Exposed NHIs in misconfigured cloud storage enabled large-scale extortion campaigns in August 2024. Similarly, Internet Archive and Cisco experienced breaches via single unrotated API keys in October 2024.

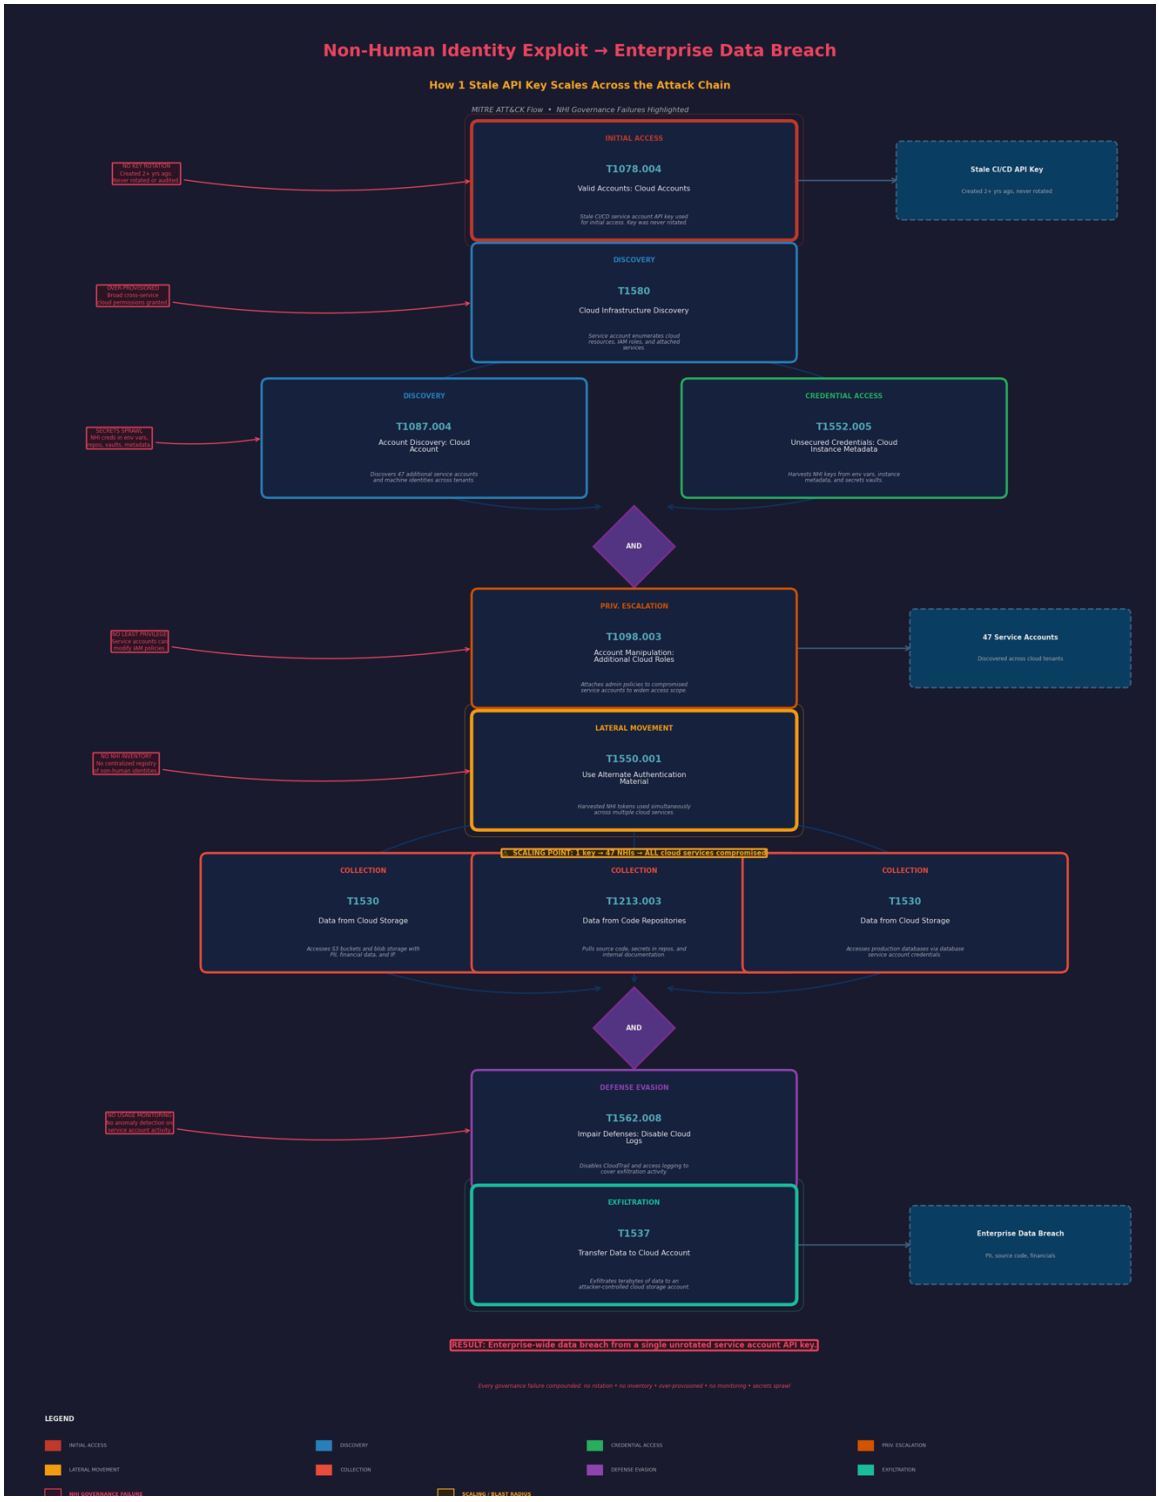

Figure 1: Attack Flow of Compromised NHI Credential Leading to Data Breach

More recent examples continue to highlight systemic risk. In December 2025, SoundCloud confirmed unauthorized access affecting 20% of its users after attackers exploited compromised identity credentials tied to SSO dashboards. Early 2026 saw Betterment suffer a breach following Okta SSO credential compromise through social engineering by the ShinyHunters group, with multiple platforms impacted by the same wave of attacks[4].

Across all these cases, the core issue remains the same: attackers did not rely on phished humans or zero-days in production infrastructure. Instead, they exploited privileged non-human identities. These identities act as doors into critical systems, and any failure in governance including excessive permissions, lack of monitoring, or stale credentials provides attackers with a trusted path into enterprise environments. The recurring lesson is clear: securing NHIs is no longer optional; it is central to preventing enterprise compromise.

Implications for Organizations

Enterprises now face an insider risk that does not align with traditional workforce identity models. Non-human insiders hold credentials, receive delegated trust, and can directly affect production systems. Unlike human employees, they operate continuously, do not require interactive sessions, and often function below the threshold of routine security monitoring. This shifts both the concentration of risk and the assumptions underlying detection and response.

The critical issue is not simply that NHIs have broad access. Most organizations still lack a unified control plane to manage them as identities rather than as discrete application artifacts. Security teams may track secrets in one system, certificates in another, cloud workload identities in a third, and AI agents nowhere durable at all. This fragmentation erodes ownership, weakens accountability, and complicates incident response. When a human account exhibits anomalous behavior, teams know how to suspend, investigate, and reauthorize it. When a non-human identity behaves unexpectedly, organizations often must first determine whether the identity exists in any authoritative inventory.

This creates a systemic risk for CISOs and enterprise risk leaders. The challenge is not only the potential for compromise, but also the presence of unmanaged authorization embedded throughout core operations. Organizations may already rely on NHIs whose full access paths are unknown. AI agents highlight this risk because they operate closer to business logic and user-facing outcomes, but the governance gaps extend across pipelines, integrations, and infrastructure. In practice, an enterprise that does not fully control its NHIs cannot fully control its business processes.

Organizational Mitigation Strategies

To mitigate the risks posed by non-human identities, we recommend the following organizational strategies:

- Build a unified NHI inventory: Every service account, API key, OAuth token, certificate, workload identity, CI/CD credential, automation bot, and AI agent must appear in one governed inventory tied to an owner, a business function, a privilege scope, a creation source, and a last-used record. Asset discovery without ownership is only enumeration. The goal is to create an authoritative identity layer for non-human actors so that every privileged machine identity has a name, a purpose, and an accountable team.
- Enforce time-to-live and automated rotation: Long-lived credentials remain one of the most common preconditions for non-human identity compromise. Every token, key, certificate, and workload credential should expire by policy and rotate automatically before that expiration window closes. Rotation must become a platform property, not an operational aspiration, because manual renewal guarantees drift and leaves stale trust paths active long after their business justification ends.
- Reduce privilege to executable necessity: Non-human identities routinely accumulate permissions because broad scopes reduce friction during deployment and troubleshooting. That convenience creates hidden insider risk. Permissions should align to the narrowest action set required for the service, workflow, or agent to complete its assigned task, and organizations should review those scopes continuously as systems evolve rather than freezing them at creation.

- Monitor runtime behavior, not just static entitlements: AI agents and other dynamic NHIs can act harmfully while remaining inside their nominal permission set. Logging successful authentication and approved API calls does not reveal whether the identity is behaving in line with its intended function. Organizations need behavioral baselines for data access patterns, tool invocation sequences, repository interactions, and downstream actions so they can detect a machine identity that begins operating like an attacker-controlled insider.

Conclusion

Non-human identities have become the backbone of modern enterprise operations, enabling automation, cloud workloads, SaaS integrations, and AI-driven workflows. At the same time, they introduce a persistent and often invisible attack surface. Recent breaches demonstrate that attackers consistently exploit over-privileged, unmonitored, or stale machine credentials to gain trusted access, bypass traditional defenses, and move rapidly through critical systems. The recurring pattern is clear: compromise rarely begins with a human user and often starts with the non-human identities that enterprises rely on every day.

Organizations that continue to treat these identities as ephemeral application artifacts rather than accountable actors face significant operational and security risk. Fragmented governance, unmanaged privileges, stale credentials, and insufficient monitoring collectively enable attackers to operate undetected. Without a deliberate approach, enterprises are delegating trust to entities they cannot fully see or control.

Mitigating these risks requires a comprehensive and unified approach to NHI governance. Enterprises must maintain authoritative inventories, enforce automated rotation, apply least privilege principles, monitor runtime behavior, and treat AI agents with the same rigor as other privileged identities. When these practices are implemented, non-human identities become governed, observable, and accountable components of the digital infrastructure rather than hidden vulnerabilities.

Securing NHIs is no longer optional. It is central to protecting critical systems, preserving trust, and maintaining operational continuity. Organizations that act now can close the governance gap, reduce exposure, and ensure that the identities driving automation remain assets rather than liabilities. A follow-up article will focus on AI agents as the emerging face of non-human identities and examine their unique risks and governance requirements.

Further Reading

1. https://www.obsidiansecurity.com/blog/what-are-non-human-identities-nhi-security-guide
2. https://www.cyberark.com/what-is/non-human-identity/
3. https://www.obsidiansecurity.com/blog/behind-the-breach-shinyhunters-2026-voice-phishing-campaign
4. https://cloud.google.com/blog/topics/threat-intelligence/expansion-shinyhunters-saas-data-theft
5. https://cyberscoop.com/shinyhunters-voice-phishing-sso-okta-mfa-bypass-data-theft/
6. https://hackread.com/gitguardian-reports-an-81-surge-of-ai-service-leaks-as-29m-secrets-hit-public-github/
7. https://thehackernews.com/2026/03/glassworm-attack-uses-stolen-github.html
8. https://cybersecuritynews.com/malicious-npm-packages-deliver-pylangghost-rat/
9. https://nhimg.org/40-non-human-identity-breaches
10. https://astrix.security/learn/blog/okta-breach-leaked-service-account/
11. https://www.cybersecuritydive.com/news/cloudflare-follow-on-attack-okta/706450/
12. https://astrix.security/learn/blog/breach-analysis-cloudflare-falls-victim-to-okta-attack/
13. https://astrix.security/learn/blog/11-attacks-in-13-months-the-new-generation-of-supply-chain-attacks/
14. https://www.csoonline.com/article/4125156/why-non-human-identities-are-your-biggest-security-blind-spot-in-2026.html
15. https://blog.gitguardian.com/the-state-of-secrets-sprawl-2026/
16. https://www.cpomagazine.com/cyber-security/hacking-campaign-scanned-exposed-git-config-files-made-off-with-15000-stolen-credentials/