Commvault

eBOOK

THE NON-HUMAN IDENTITY CRISIS: BRIDGING THE BLIND SPOT IN MODERN DATA PROTECTION & CYBER RESILIENCE

EXECUTIVE SUMMARY

In the era of AI-driven automation, the traditional perimeter has dissolved, leaving **identity** as the final control plane. While enterprise security has spent decades hardening human authentication, a silent explosion of **Non-Human Identities (NHI)** – service accounts, API keys, OAuth tokens, and AI agents – has created a massive, unmanaged attack surface.

These identities now outnumber humans by a ratio of **144**:1[1], yet they are rarely governed with the same rigor. This ebook explores how modern adversaries exploit this "identity debt" to achieve persistence and why organizations must pivot toward a Tier 0 recovery-first architecture to enable identity resilience.

1 The NHI & Secrets Risk Report H1 2025, Entro

SECTION 1: DEFINING THE MACHINE ATTACK SURFACE

The scope of NHIs extends far beyond cloud-native software; it represents the critical connective tissue for the physical and medical infrastructure that sustains modern society. In these high-stakes environments, an NHI is any device or "thing" that possesses credentials, executes actions, and moves data autonomously.

It is vital to recognize that the NHI "blind spot" now includes hardware assets that interact directly with the physical world. When these identities are overprivileged or under-governed, the risk involved is not just data loss but also includes physical and life-safety hazards.

FINANCIAL INFRASTRUCTURE

Modern ATMs function as sophisticated NHIs that authenticate to core banking systems via service accounts and API keys to authorize cash dispersal and process sensitive personally identifiable information.

MANUFACTURING & INDUSTRIAL IOT

In smart factories, programmable logic controllers and robotic arms act as workload identities within DevOps-style automation pipelines. A compromise here allows an attacker to manipulate integrity – altering physical production parameters without triggering human alerts.

AGENTIC AI IN PHYSICAL SYSTEMS

As agentic AI is deployed to manage building scripts or power grids, these agents authenticate to act. Malicious manipulation through prompt injection could trick an agent into exfiltrating sensitive site schematics or disrupting utility services.

BIOTECH & MEDICAL IOT:

In-hospital devices

Imaging machines and infusion pumps function as NHIs that inherit access to electronic health records, often using long-lived tokens that are rarely rotated.

Implantable devices

Connected pacemakers and insulin pumps represent the most sensitive edge of the identity plane. These devices authenticate to provider networks to upload telemetry. If the identity layer is compromised, the “business decision” being manipulated is a life-critical medical adjustment.

Identity resilience for these devices requires a **recovery-first** architecture, where the organization must be able to audit and rollback unauthorized privilege escalations in real time to verify that your physical and medical “identities” remain in a trusted state.

THE RISK IN NUMBERS

The scale of the NHI landscape is staggering and growing at a rate **4 to 10** times faster than human accounts.[2]

GOVERNANCE GAP

Fewer than

25%

of organizations have formal policies for the creation or decommissioning of these identities.[3]

EXCESSIVE PRIVILEGE

97%

of NHIs possess permissions far beyond what is required for their functional role.[4]

VULNERABILITY

Only

12%

of organizations express high confidence in their ability to prevent an attack targeting these machine identities.[5]

2 CyberRisk Alliance
3 The State of Non-Human Identity and AI Security, Cloud Security Alliance
4 2025 State of Non-Human Identities and Secrets in Cybersecurity, Entro
5 The State of Non-Human Identity and AI Security, Cloud Security Alliance

SECTION 2: CASE STUDY IN MODERN WARFARE— THE SLH AND SHINYHUNTERS MODEL

Adversaries like the cybercriminal supergroup **Scattered LAPSUS$ Hunters (SLH)** – which includes notorious actors like **ShinyHunters** and **Scattered Spider** – have pioneered a lethal shift in attack methodology. Their tactics illustrate why NHIs are more strategically valuable than human credentials.

THE INDUSTRIALIZED VISHING PIPELINE

SLH has industrialized **Voice Phishing (vishing)**, a vector that saw a **449%** increase in 2025.[6]

01 RECRUITMENT AND INCENTIVE

The group has been observed offering financial incentives of $500 to $1,000 per call to recruit specialized talent, specifically seeking women to conduct vishing attacks.[7] This tactic is designed to increase the success rate of help desk impersonation by bypassing traditional "attacker" profiles that staff may be trained to identify.

02 THE ENTRY

Using pre-written scripts, these recruits impersonate employees to convince IT help desks to perform password or multi-factor authentication (MFA) resets.

03 THE PIVOT TO NHIS

Once initial access is obtained, attackers move laterally to virtualized and cloud environments. They often maintain persistence by migrating from the human account to machine-based layers, such as stealing OAuth tokens or creating new administrative service accounts.

04 THE PERSISTENCE

Unlike human passwords, these machine identities are rarely rotated and often invisible to traditional monitoring. This allows the attacker to maintain access long after the compromised human user has been remediated.

6 2025 Phishing Threat Trends Report, Vol. 6, KnowBe4
7 The Hacker News

SECTION 3: THE ⚠ CRITICAL VULNERABILITIES OF NHIS

NHIs bypass the three traditional pillars of security: user MFA, endpoint detection, and email filtering.

OAUTH ABUSE

Attackers trick users into approving applications through a consent screen, granting delegated access to mailboxes and data. This activity appears as normal API traffic and remains active even after password resets, which makes it harder to acknowledge and register this as a threat to warrant any action.

SERVICE ACCOUNT GHOSTING

In large enterprises, thousands of undocumented service accounts operate with administrative privileges, often remaining active long after their original project has ended. This lack of hygiene is a key area that needs to be addressed effectively and efficiently.

TOKEN & KEY THEFT

Long-lived API keys embedded in static configuration files or DevOps pipelines lack device or geolocation context, granting broad, automated access without triggering interactive login alert.

SECTION 4: A FRAMEWORK FOR IDENTITY RESILIENCE

Organizations must evolve their workflows to treat NHIs as **Tier 0** assets, equivalent in risk to domain administrators or cloud control planes.

PRIVILEGED CREATION PROTOCOLS

The creation of an OAuth app or service account should require administrative approval and generate high-signal security telemetry.

HUMAN-TO-NHI CORRELATION

Most breaches could be stopped early by correlating cross-domain signals, such as a help desk interaction followed immediately by an MFA reset or a new token creation.

ELIMINATION OF STATIC SECRETS

Static secrets represent an unacceptable risk. They must be replaced with short-lived tokens and automatic rotation.

RECOVERY-FIRST ARCHITECTURE

Since 100% prevention is impossible, organizations must be able to quickly detect, contain, and recover the identity environment to a trusted state.

SECTION 5: COMMVAULT IDENTITY RESILIENCE

Commvault identity resilience helps provide the visibility and recovery capabilities required to protect complex identity environments at enterprise scale.

FAST IDENTITY RECOVERY

Provides automated recovery capabilities ranging from individual AD attributes and objects to complete forest restoration. Because AD forest recovery is one of the most complex and error-prone recovery processes in IT, involving DNS, SYSVOL, FSMO roles, Global Catalogs, trusts, replication, and recovery sequencing across multiple domains, the solution is designed to automate and orchestrate these steps through guided workflows and runbooks.

This helps organizations rapidly restore the identity environment to a known-good, trusted state, minimize downtime, reduce the risk of human error, and reestablish the authentication and authorization foundation on which the rest of the organization depends.

VULNERABILITY ASSESSMENT

Evaluates the Active Directory (AD) environment for security weaknesses that increase the likelihood or impact of compromise. The assessment identifies excessive privileges and over-permissioned accounts, use of legacy and insecure protocols such as NTLMv1, SMBv1, and unsigned LDAP, weak authentication settings, missing hardening controls, and risky configuration drift.

It also highlights issues such as unconstrained delegation, stale privileged accounts, weak password and Kerberos policies, insecure trust relationships, and lack of protections for Tier 0 assets. The result is a prioritized map of identity risk, helping organizations reduce the blast radius and strengthen identity resilience before an attacker can transform AD into a persistent platform for compromise and lateral movement.

REAL-TIME AUDITING & ROLLBACK

Continuously monitors AD using directory synchronization data and Windows event logs to detect suspicious or unauthorized changes as they occur. The service tracks activities such as privilege escalation, additions to privileged groups, changes to GPOs, modification of Tier 0 accounts, authentication policy changes, creation of rogue users, and attempts to weaken security settings.

Each event is correlated into a timeline so administrators can quickly distinguish routine administration from the opening moves of an identity attack. When malicious or accidental changes are identified, the platform is designed to support rapid, one-click rollback to help restore the previous trusted state.

Learn more about identity resilience

commvault.com | 888.746.3849 | get-info@commvault.com