Silverfort

Insecurity in the shadows:
New data on the hidden risks of non-human identities

Shining the spotlight on NHIs so you can find, fix and fortify them before they're exploited.

NHIs are a force multiplier—for you and your adversaries

If you've glanced at any security blog or news outlet in the past few months, you've probably heard about the rising risks of non-human identities (aka NHIs).

NHIs is a new(ish) term used to describe when a machine, application or service is given credentials to perform an automated task or action. In the on-prem world, the most common NHIs are service accounts, but there are many other types too, including API keys, system accounts, and OAuth tokens. They're all essential to the functioning of modern environments both in the cloud and on-prem—and have been for a long time.

But there's a problem: adversaries don't just hack people. NHIs are notoriously under-observed, under-protected and over-privileged—and their numbers are growing exponentially. Combine these four factors together, and it's easy to see why NHIs are prime targets for attackers seeking ways to slip through the cracks and move undetected through environments.

We're here to cut through the complexity and get to what really matters: what the NHI landscape looks like today, what these risks mean for you, and what you can do about them.

Why are NHIs a major security risk?

Non-human identities: The silent workers who hold the keys to the kingdom.

They're under-observed

NHIs are easy to create and difficult to monitor. With no centralized visibility, no organized onboarding and offboarding process, and a chronic lack of ownership, NHIs are often left to their own devices.

5.7% **of organizations have full visibility into their on-prem service accounts.**

They're over-privileged

NHIs typically have high access privileges in excess of what they need to complete their tasks. If compromised, this can allow access outside of their intended use, making them lucrative targets for lateral movement.

35% **of all users are service accounts with high access privileges and low visibility.**

They're under-protected

NHIs cannot be protected like human users. Multi-factor authentication (MFA) is not applicable, and credentials/secrets cannot be easily rotated in a privileged access management (PAM) vault due to the risk of crashing critical processes.

80% **of organizations cannot prevent the misuse of service accounts in real time.**

They're everywhere

We estimate that NHIs now outnumber human users 50 to 1—and that number will continue to grow. The scale of the problem—and the growing attack surface that comes with it—is rapidly becoming unmanageable.

NHIs outnumber human users 50 to 1 in 2025.

THEY'RE UNDER-OBSERVED

NHIs are easy to create but difficult to monitor

Most organizations don't have full visibility into their NHIs. In fact, only 5.7% of security leaders believe they know where all their NHIs are and what they do, making it hard to detect misuse and even harder to secure every single one of them. Within the large pool of NHI types, service accounts—used for machine-to-machine communication within Microsoft's Active Directory's (AD) environments—are very concerning. Our data shows that on average 35% of a company's user accounts are service accounts.

35%

of a company's user accounts are service accounts.

It began with human admins creating NHIs across multiple systems, apps, platforms and services without a centralized view of what they are responsible for, who created and owns them, and what they need access to. For most organizations, there has never been a single source of truth for NHIs, nor any standardized onboarding, offboarding, or ownership processes, leaving their inventories incomplete at best.

The shift to the cloud exacerbated this issue—and set the stage for Gen AI, Large Language Models (LLMs) and Copilot to autonomously and silently use NHIs.

Without ongoing visibility into what is being created and why, organizations are forced to play catchup to their growing identity attack surface. After all, you can't protect what you can't see.

Level of visibility into service accounts

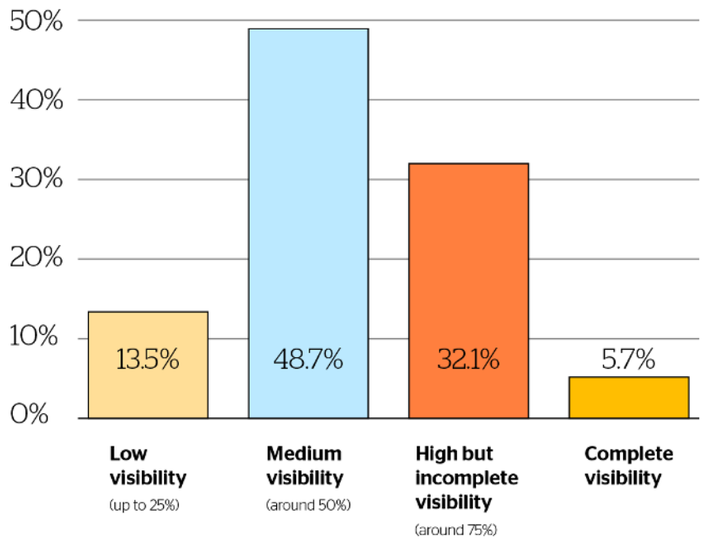

**Osterman State of Identity Attack Surface Report 2023*

THEY'RE UNDER-PROTECTED

NHIs cannot be protected like human users

Traditional identity security controls are almost entirely human-centric. And securing your NHIs is not the same as securing your privileged human users. Even more, New Technology LAN Manager (NTLM) still exists in many Windows domains despite being a very weak authentication protocol that's susceptible to credential access and lateral movement. In fact, 46% of service accounts regularly authenticate via this deprecated protocol, leaving them more exposed to compromise.

NHIs can't use MFA like humans can, and the risk of breaking vital processes by protecting them in a PAM vault with password rotation often outweighs the benefits. Similarly, as we've already seen, NHIs are not typically subjected to same level of scrutiny as human users, with formalized onboarding and offboarding procedures being a rare occurrence.

This security challenge is reflected in how confident organizations feel in preventing the misuse of their NHIs. Four in five organizations do not trust that they can prevent adversaries from using an NHI for malicious access due to sporadic or absent visibility and security. Coupled with the fact that 80% of organizations have experienced an identity-related breach, this paints an alarming picture.

46% **of service accounts regularly authenticate with NTLM, a deprecated protocol.**

37% **of on-prem service accounts are "interactive".**

These are service accounts that appear to be used by employees to bypass privileged access controls or by an attacker.

80%

of organizations cannot prevent the misuse of service accounts in real time.

THEY'RE OVER-PRIVILEGED

NHIs have high privileges—often more than they need

NHIs play a vital role in keeping organizations running smoothly. As such, they often require access to specific resources and permissions tailored to their functions, which can lead to broader access scopes than those typically assigned to human users.

But the issue goes deeper, as we commonly see NHIs with excessive privileges. When establishing machine-to-machine access for NHIs, it's often far easier and more convenient to grant broad access rather than fine-tune their permissions and ring-fence them to their specific purpose. The risk of friction is lower—but the potential cost to security is much higher.

NHIs' high access privileges, coupled with the fact that they're under-observed and under-protected, make them lucrative targets for attackers seeking ways to perform lateral movement and privilege escalation. 56% of organizations unknowingly sync more than half of their service accounts to their SaaS directory.

Here's an example of how a breached NHI (in this case, an AD service account) can also lead to the compromise of an organization's entire SaaS environment:

Even though service accounts are not supposed to be synced from AD to the cloud identity provider (IdP), it's extremely common for identity teams to sync them inadvertently. While these accounts can't be used to access SaaS resources by default, an attacker that has gained admin access privileges to the cloud IdP can activate them and assign them access privileges.

Sample attack flow

1. Attacker uses credentials obtained either in a recent third-party supplier breach or from phishing the credentials of the IT administrator using a fake login page.
2. They use this access to get to PII data stored in a Snowflake DB (or any other SaaS app from the IdP access they have).
3. They then search for synced service accounts (naming conventions are a useful guide) until finding one to gain access to the on-prem environment.
4. Once inside the on-prem environment, the attacker can run code, exfiltrate data, or continue to move laterally.

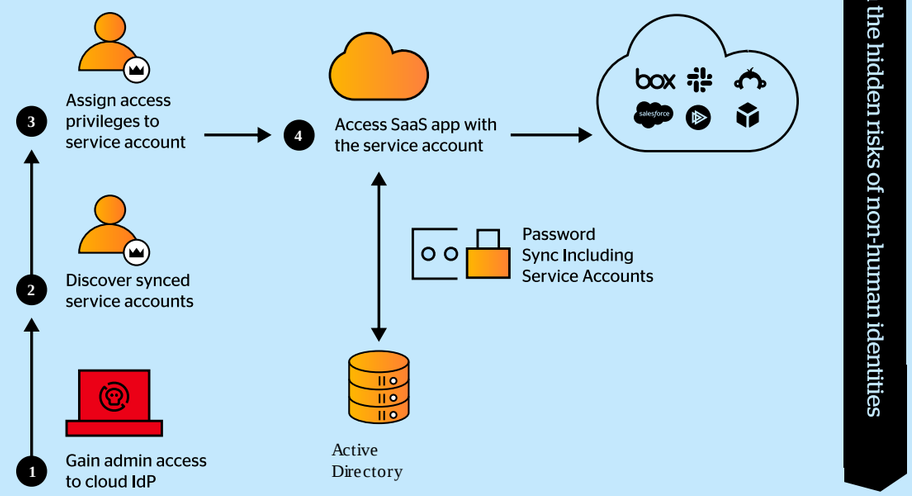

THEY'RE EVERYWHERE

The volume of NHIs is growing—fast

NHIs already outnumber human identities in most hybrid environments. As cloud, SaaS and API adoption grows, so do these identities—often without identity and security teams knowing where they are and what they're doing.

They now make up a very significant portion of total users. In 2020, NHIs outnumbered human users by 10 to 1. It's now estimated to be more like 50 to 1. In hybrid environments we found a ratio to between 1:30-1:50 for human to NHIs, while in multi-cloud environments the numbers are even higher. IAM roles, secrets and service accounts are the most common types of NHI in the wild, with a major lead over other types on NHI.

Almost in all cases, customers report hardships in assigning NHI's human owners and therefore lacking the ability to protect NHIs end-to-end or rapidly remediate posture issues. More than half of customers estimate that 40% of NHIs have an unknown owner, many of whom are dormant or set and forget.

While the sheer number of NHIs in the typical enterprise environment isn't a problem in itself, it does add a layer of complexity and urgency to the three issues we've already discussed: visibility, protection, and privileges. As this volume continues to multiply, so does the task of discovering, monitoring and securing every NHI, leaving organizations forever a step behind their rapidly expanding identity attack surface.

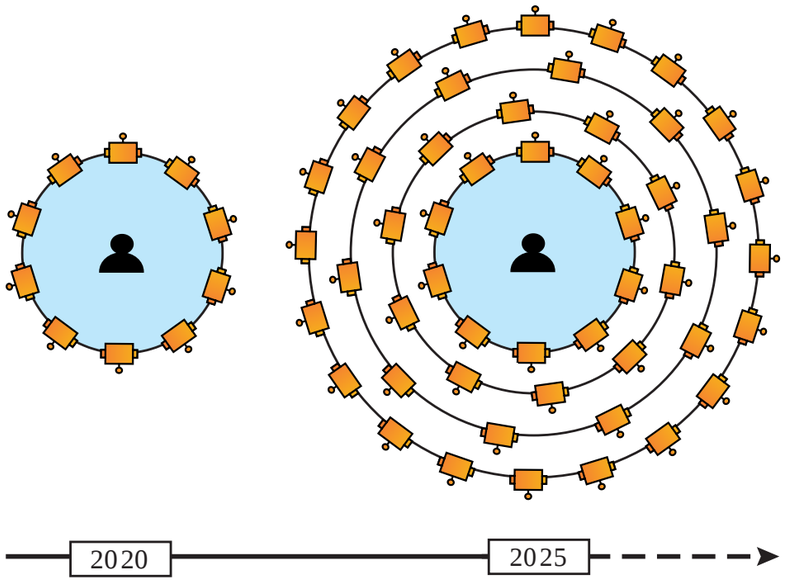

40%

of NHIs have an unknown owner

The role of NHIs in real-world attacks

Attackers will target NHIs for lateral movement due to their high-access privileges, low visibility, and protection challenges. Even better—in many cases, NHIs fly under the radar of security and identity teams because they don't even know they exist.

1 API key or token theft

Internet Archive Breach (October 2024):
Attackers exploited unrotated API keys leaked from the Internet Archive's GitLab repository, gaining access to over 800,000 support tickets containing sensitive user information.

2 Overprivileged service accounts

Dropbox Sign Breach (May 2024):
Attackers compromised a backend service account with excessive privileges, accessing the customer database and exposing sensitive user data, including email addresses, usernames, hashed passwords, API keys, and OAuth tokens.

3 OAuth application abuse

Microsoft and Okta Attacks:
Nation-state actors have been seen to abuse OAuth applications to move laterally across cloud environments. Major software companies like Microsoft and Okta have fallen victim to attacks leveraging compromised machine identities, highlighting the vital need to connect non-human identities with their human counterparts for complete visibility and protection.

How to secure your non-human identities

Applying the IDEAL Framework to NHI Security

Yesterday's solutions cannot solve today's problems. As the need for effective NHI security grows more urgent—and the sheer scale of the problem continues to escalate—it's time for businesses to demand more of their identity security.

To systematically address NHI security challenges, we suggest taking the IDEAL identity security approach. If you want an in-depth look at this approach, check out the expertly authored Identity Security Playbook.

Get the Identity Security Playbook

Integrate with all IdPs: on-prem AND cloud

Connect to all identity providers—on-prem and cloud—to capture authentication and access attempts, ensuring comprehensive monitoring of NHIs across your entire hybrid environment.

Discover and classify all NHIs

You can't protect what you can't see. Identify all NHIs across cloud workloads, SaaS apps, AD and all other IdPs and categorize them based on risk and function. Connect every NHI to their human counterparts for full visibility and lifecycle management.

Enforce security controls

While NHIs cannot be protected in the same way as human users (with, for example, MFA), virtual fencing can be used to block suspected malicious access, and eliminating excess privileges can make it much harder for attackers to use them for nefarious means.

Analyze and act on all NHI activity

NHIs tend to have a very clear, predictable pattern of behavior. Mapping this behavior, along with their destinations, sources, privilege levels and security posture will establish a baseline. Flag any deviations from this baseline and implement automated, real-time responses.

Lightweight deployment and maintenance

Deploy solutions that enhance NHI security without causing disruptions or adding to your team's workload. Solutions that can continuously and automatically update your NHI inventory will make the biggest impact on your overall NHI security posture.

Silverfort Identity Security Platform

NHIs, whether on-prem or cloud, are the backbone of modern IT. They're also one of the biggest security gaps. **If you're not actively securing them, attackers will exploit them.**

The good news? By taking **a proactive, automated approach**, you can dramatically reduce your exposure.

And even better: we've got just the thing to help.

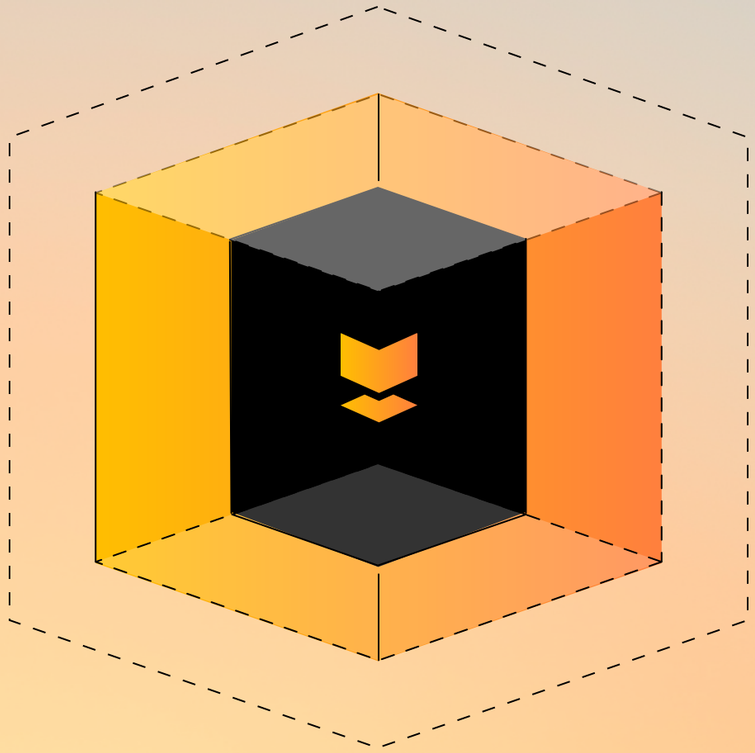

See it in action

Watch a video of Silverfort's non-human identity security, or try it hands on with our interactive tour.

Product tour

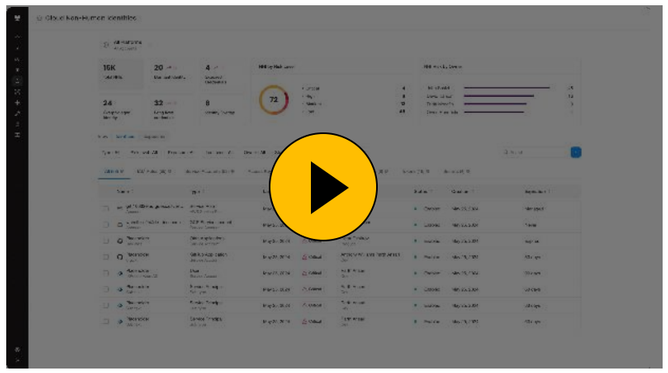

Video

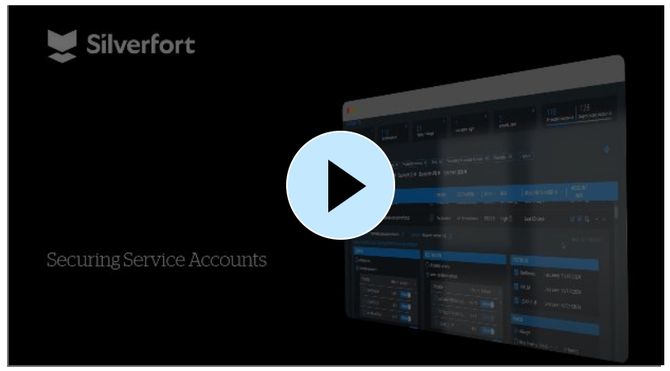

Report methodology

In Silverfort's Identity Underground Report, we looked at hundreds of thousands of data points across hundreds of customers of different sizes and verticals to determine the scope of the non-human identity problem, with a focus on highly privileged and pervasive Active Directory service accounts.

We also conducted research with Osterman Research, which included responses from 637 people in identity roles during May-June 2023. To qualify, respondents had to work at organizations with at least 1,000 employees. The surveys were conducted in six countries, with the surveys in France and Germany fielded in French and German respectively. The survey was cross industry, and no industries were excluded or restricted.

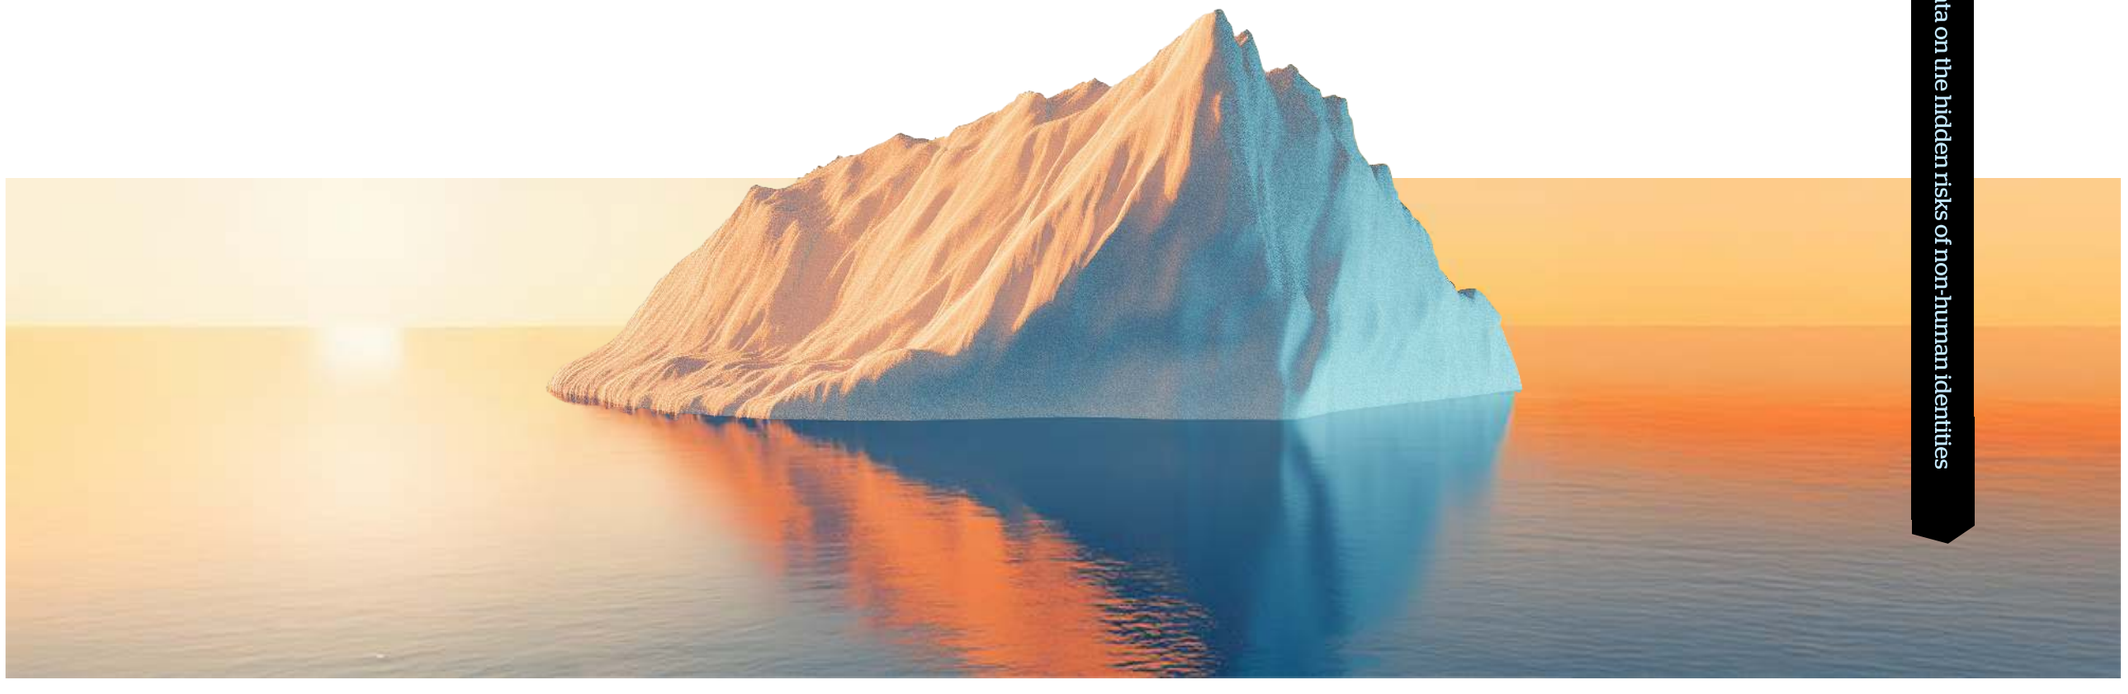

About Silverfort

Silverfort secures every dimension of identity. We are the first to deliver end-to-end identity security across the entire IAM infrastructure, eliminating gaps and blind spots, giving organizations visibility into their identity fabric and extending protection to resources that previously could not be protected. This is all done via a patented technology that natively integrates with the entire IAM infrastructure, Runtime Access Protection™ (RAP). It is lightweight, easy to use and deploy, and won't disrupt business operations, resulting in better security outcomes with less work.

Discover every identity across every environment, **analyze** exposures to reduce your attack surface, and **enforce** security controls inline to stop lateral movement, ransomware propagation, and other identity threats.

To learn more about Silverfort, **visit silverfort.com**