Token Security

THE CISO'S COMPLETE GUIDE

TO AGENTIC AI AND NON-HUMAN IDENTITY SECURITY

token

Table of Contents

With the rapid rise of agentic AI, non-human identities (NHIs) have quietly become the largest, least-governed attack surface in the modern enterprise. What was once considered a low-priority mitigation afterthought is now a board-level business risk.

This guide provides CISOs with a comprehensive roadmap for tackling the complex security challenges of NHIs, which includes agentic AI, where autonomous agents can expose secrets using stale credentials at machine speed.

token

01

AGENTIC AI AND NON-HUMAN IDENTITY CHAOS:

THE NEW HIDDEN ATTACK SURFACE

There is a hidden threat plaguing organizations today. That threat? AI agents and non-human identities (NHI).

Non-human identities (NHI) are digital, programmable credentials that interact with applications or systems to automate processes, such as accessing cloud environments, API authentication, AI infrastructure, AI frameworks, and more, without requiring human intervention. NHIs include OAuth tokens, API keys, secrets, service accounts, workloads, AI agents, and more.

Non-human identities outnumber humans by a ratio greater than 45:1

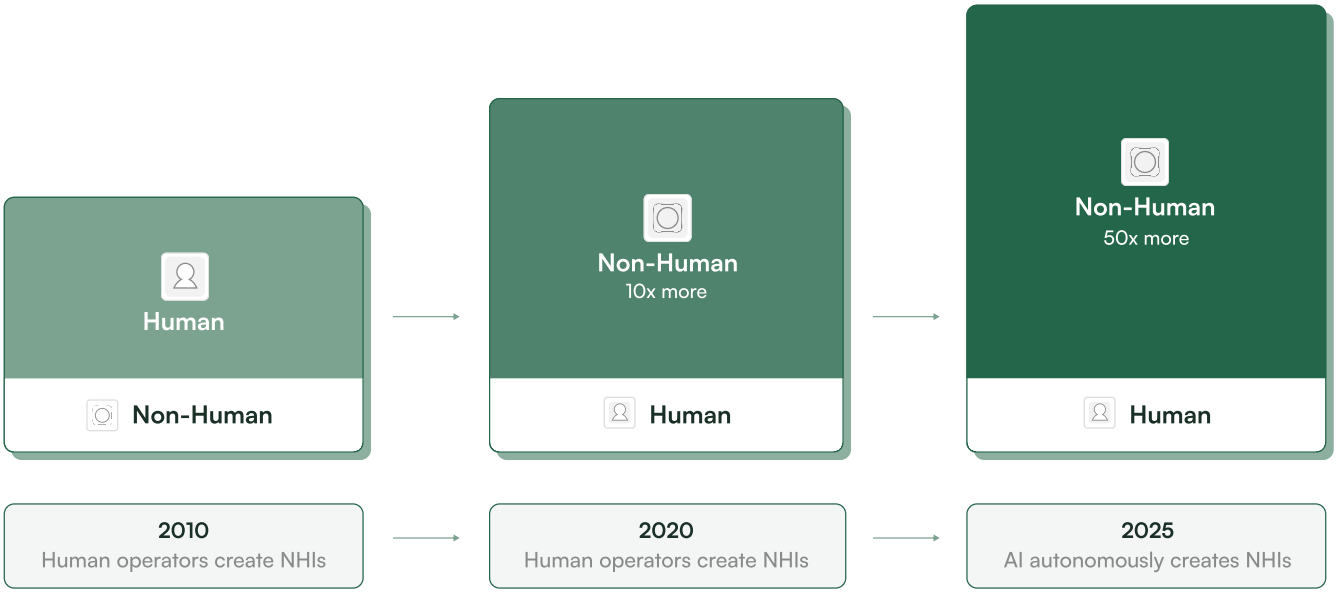

NHIs are decentralized across both cloud and on-premises environments, making them inherently difficult to monitor, govern, manage, and secure. Without clear ownership and accountability, CISOs cannot enforce the right access and security policies, adding to the growing identity complexity attack surface and compounding the risks for an NHI-related breach. But this is where the identity chaos begins.

Two-thirds of enterprises have experienced a successful cyberattack resulting from compromised non-human identities

NHIs also present limited visibility into AI services, SaaS applications, CI/CD pipelines, microservers, containers, and third parties, where identities are often provisioned outside of the organization's direct control. These critical blind spots make it increasingly difficult for CISOs to revoke excessive permissions, remove orphaned accounts, and track identity lifecycle management.

token

02

THE CHALLENGES OF MANAGING AI AGENT AND NON-HUMAN IDENTITIES

AI agents and non-human identities introduce identity-related risks at a volume and speed that traditional identity and access management (IAM) tools weren't designed to handle. This not only makes them ineffective for managing modern identity sprawl, but it also places additional pressure on CISOs when justifying cybersecurity investments to the board. Furthermore, these legacy IAM tools fail to provide the insights required to govern machine-to-machine interactions at scale.

A Ponemon study found that only 50% of organizations rate their IAM tools as very or highly effective

Traditional IAM tools fall short in managing AI agents and NHIs across several key areas:

Discovery and Visibility:
Legacy IAM tools are built to detect and manage human users in structured directories. They lag in discovering AI agents and NHIs across modern environments like cloud platforms, CI/CD pipelines, containers, and AI environments, which go unmanaged for long durations of time.

Complex Identity Lifecycle Management:
Without automation, provisioning and deprovisioning AI agents and non-human identities, rotating secrets and API keys, reviewing permissions, and enforcing least privilege policies becomes nearly unmanageable.

Lack of Visibility:
Traditional IAM tools lack automation and visibility, leading to orphaned accounts, unrotated credentials, expired tokens still in use, over-permissioned service accounts and agents, and hardcoded secrets in application code. Since there is no centralized tracking, there is no concrete way to determine which are active, who owns them, and what level of access they have.

Ownership:

Without clear attribution, AI agents and non-human identities often operate independently without an accountable owner. Traditional IAM tools lack the contextual metadata to link AI agents and NHIs to specific individuals or business-related functions. Decommissioning agents and identities becomes exceptionally challenging for security teams since there is no defined owner responsible for ensuring their lifecycle is properly managed or tracked. Access reviews become ineffective, as there is no simple method to identify who or what is actually using the identity or what permissions are attached.

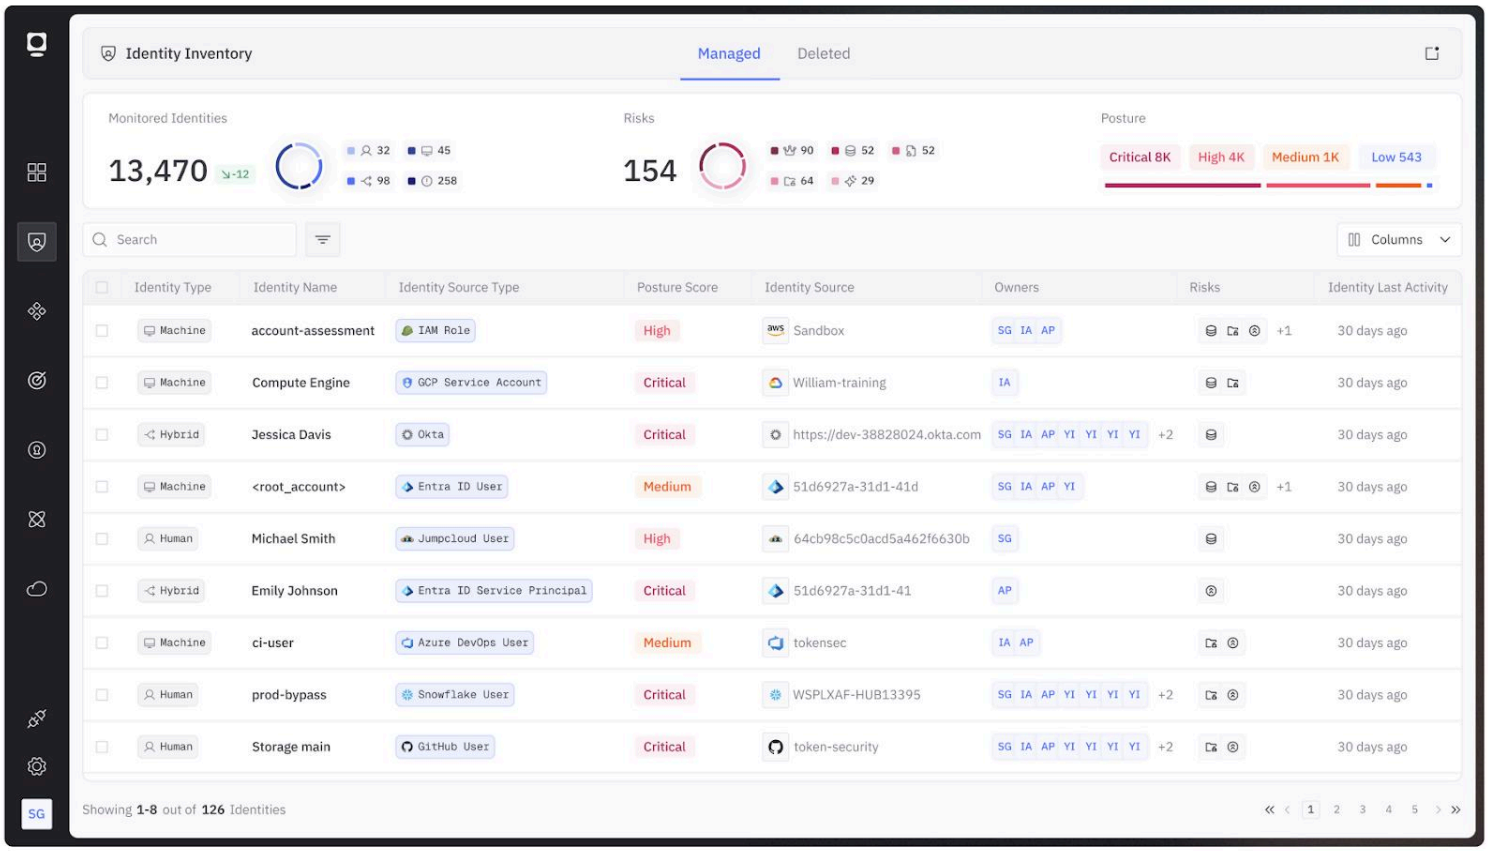

token

03

CRITICAL SECURITY RISKS OF AI AGENTS AND NON-HUMAN IDENTITIES

66% of organizations experienced a security incident caused by a compromised NHI

AI agents and non-human identities present inherent risks that, if left undetected or unmitigated, can lead to significant identity-related security breaches.

Here are the critical security challenges:

39 million leaked secrets were detected on GitHub in 2024

35% of exposed API keys are still active

1. Unrotated Credentials

Credentials can linger for extended periods in service accounts, cloud environments, and AI services without being updated or revoked. Leaked secrets, API keys, or OAuth tokens hardcoded in a script or left undetected in a GitHub repository can lead to unauthorized access and privilege escalation.

Prevention

- Enforce regular rotation policies for all AI agent and NHI credentials
- Use temporary credentials and set expiration policies
- Leverage automation tools to rotate API keys, service account keys, and other credentials

31% of organizations have former employees who have accessed assets stored in SaaS applications

2. Incomplete Offboarding

What happens when a lead DevOps engineer leaves the organization? That engineer might have hundreds of unknown NHIs across fragmented developer environments, CI/CD pipelines, Kubernetes clusters, GitHub repositories, IAM roles, and AWS accounts. They might have secrets embedded in Helm charts or container images stored in artifact registries. Additionally, there may be orphaned AI agents that still connect to critical services. Former employees might also have access to sensitive company data through SaaS applications using OAuth tokens that were never revoked.

Prevention

- Establish clear offboarding policies for AI agents, NHIs, and accounts no longer in use
- Audit and inventory all AI agents and NHIs across all environments (SaaS integrations, cloud accounts)
- Automate secret rotation

23% of cloud identities have critical or high-severity excessive permissions

3. Overprivileged Identities

Can you guess how many NHIs have access to AWS, Azure, or GCP environments? What about AI agents that have access to business critical applications and sensitive data. This presents many blind spots, particularly with permission sets. Large-scale enterprises might have a finite number of EC2 workloads with excessive read and write permissions that have access to sensitive data stored in S3 buckets.

31% of S3 buckets are currently open to the public

When was the last time you reviewed permission scopes across your cloud accounts? A misconfigured S3 bucket can leak sensitive data and trigger a cloud-related breach.

Prevention

- Restrict access permissions by implementing the principle of least privilege
- Continuously monitor and alert on permission changes
- Conduct periodic audits of permissions given to AI agents, service accounts, and NHIs

65% of organizations have not implemented MFA in a way that provides comprehensive security

4. Weak Authentication

Weak Authentication: Hardcoded passwords in a Python script can lead to credential exposure if the code is shared or deployed in environments. AI agents and NHIs authenticating autonomously using these static credentials often bypass standard identity guardrails, leading to brute force attacks, credential stuffing, and lateral movement. Service accounts are at high risk, especially given the volume of third-party NHIs with overpermissions or unrestricted access

Credential stuffing attacks cause on average $4.81 million worth of damage per breach

Prevention

- Use SSH keys and OAuth tokens instead of traditional passwords whenever possible
- Ensure that MFA is enabled across all accounts
- Enforce strong password policies with regular updates

> **A non-human identity (NHI) solution is quickly becoming an essential tool for modern IT and Engineering environments and a crucial part of every comprehensive security program. Token has built a very strong team and a compelling product that knows how to address these problems and is a critical tool to implement for covering NHI risks spanning from legacy applications to modern AI agents.**
>
> **Chad Kalmes,** CISO at Udemy

token

04

OWASP TOP 10 NON-HUMAN IDENTITIES RISKS FOR 2025

Improper Offboarding tops the list of non-human identities risks. This is particularly relevant when offboarding employees, as their NHIs can live on in the form of orphaned accounts, OAuth tokens, and API keys that remain active long after the employee has left the organization. AI agents can trigger automated workflows endlessly and undetected, even using stale credentials or tokens.

What should you do?

Below is the OWASP Top 10 Non-Human Identities Risks for 2025 table with best practices on prevention.

Improper Offboarding

Best practices:

- Discover every NHI and reassign ownership
- Rotate and revoke credentials
- Regularly audit active NHIs

Secret Leakage

Best practices:

- Automate secrets detection in CI/CD pipelines and code repositories
- Use short-lived credentials whenever possible
- Implement the principle of least privilege for access permissions

Vulnerable Third-Party NHI

Best practices:

- Conduct comprehensive security reviews before integrating third-party tools
- Automate credential rotation
- Continuously monitor third-party NHI activity

Insecure Authentication

Best practices:

- Use OAuth 2.1 and OIDC for secure authentication
- Perform regular authentication reviews and remove unused accounts
- Replace static credentials with temporary, scoped tokens

Overprivileged NHI

Best practices:

- Regularly audit and review permissions
- Implement Just-in-Time (JIT) access for elevated privileges
- Enforce the principle of least privilege access for each NHI

Insecure Cloud Deployment Configurations

Best practices:

- Conduct periodic cloud configuration reviews
- Restrict trust relationships in IAM roles and OIDC configurations
- Perform Infrastructure-as-Code (IaC) scanning to detect any misconfigurations before deployment

Long-Lived Secrets

Best practices:

- Implement automated rotation policies for API keys, tokens, and credentials
- Regularly check code repositories and CI/CD pipelines for leaked secrets
- Enforce the principle of least privilege

Environment Isolation

Best practices:

- Enforce environment-specific access controls
- Apply the principle of least privilege for each environment
- Assign environment-specific NHIs

NHI Reuse

Best practices:

- Assign environment-specific NHIs
- Assign unique NHIs to each application or service
- Conduct regular NHI audits

Human Use of NHI

Best practices:

- Ensure that NHIs are provisioned automatically
- Audit and monitor NHI activity
- Educate developers and administrators on the risks of using NHIs

OWASP Non-Human Identities Top 10

token

05

THE RISE OF AGENTIC AI:

THE NEXT FRONTIER IN NON-HUMAN IDENTITY SECURITY RISK

AI agents add a new level of complexity to the non-human identity chaos landscape. AI agents rely on non-human identities to authenticate, access systems, and communicate with APIs. Gartner predicts that "By 2027, AI agents will reduce the time it takes to exploit account exposures by 50%." That's the good news.

Nearly 12,000 API keys and passwords were found in an AI training dataset

However, the challenge with non-human identities is that they often operate autonomously, without human oversight. LLMs cannot differentiate secrets during training, leaving sensitive credentials and API keys at risk of exposure if they're embedded in training datasets.

An attacker could essentially manipulate the AI agent with a prompt injection to escalate privileges further or exploit access to the machine learning model. Factor in third parties, and the complexity of identity management grows exponentially.

There are several considerations a CISO should take into account when securing AI agents with non-human identities

- **What permissions do the AI-created NHI agents have?**
- **Do the AI agents have least privilege access enabled?**
- **Is there a clear process for decommissioning AI agents if they are no longer needed?**
- **Is there a regular review of permissions and roles assigned to AI agents?**
- **How is the training data for AI agents secured to prevent unauthorized access?**
- **Are expiration policies applied to temporary AI-created NHIs?**

token

06

WHY CISOS NEED TO REEVALUATE IDENTITY MANAGEMENT IN 2025

The global machine identity management market, valued at $1.2B in 2021, is projected to reach $2.4B by 2027

The dependence on traditional IAM tools and other fragmented authentication systems creates more complexity in terms of visibility, governance, and control over AI agents and non-human identities. One particular challenge is the risk of shadow identities, or unmanaged non-human credentials, such as AI agents, API keys, tokens, and service accounts, that are created outside approved workflows. These hidden identities reside in decentralized environments and often go undiscovered.

99% of cloud identities are overly permissive, and only 1% of the identities are granted least-privileged permissions

This is where a machine-first identity security approach is needed. The table below highlights the contrast between a human-centric approach and a machine-first identity security approach.

Point-by-Point Comparison

Human-Centric IAM	Machine-First Identity Security
Limited visibility into unmanaged NHIs	Unified identity visibility
No clear ownership or accountability of NHIs	Full attribution data tied to the owners of each identity
Prone to policy misconfigurations	Automatically enforces least-privilege access policies
Manual provisioning, deprovisioning, and credential rotation	Automates identity lifecycle management of AI agents and NHIs
Risk of overpermissions and unrotated credentials	Automatically rotates API keys, tokens, credentials, and certificates
Remediate without context	Prioritizes remediation efforts based on business risk

token

07

10 KPIS A CISO CAN PRESENT TO THE BOARD

78% of organizations expect to increase their identity security spending in 2025

AI agents and non-human identities might sound like foreign concepts to the board, and new initiatives typically require a lot of justification to secure capital expenditures. CISOs must articulate risk in business terms and objectives to overcome any communication barriers or pushback. Speak the language of the board, in terms of financial security risk exposure and ROI. The best way to do that is through business-driven outcomes, or KPIs.

Research from Splunk's CISO Report found that 64% of boards say presenting security as a business enabler is the most effective way to increase budget. No better way to do that than by demonstrating how securing AI agents and non-human identities directly supports business goals, fosters innovation, and reduces "invisible" risk.

Here are 10 business-driven KPIs that can be measured and quantified month-over-month (MoM), quarterly, or during annual board reviews.

10 Security KPIs for AI Agents and NHIs

Number of Inactive Employee Accounts Successfully Offboarded	Number of Orphaned Accounts Removed
Percentage of Accounts that Have MFA Enabled	Total Number of Privileged Accounts Reduced
Number of Hardcoded Credentials Detected	Average Time to Revoke Compromised Credentials
Number of Users with Excessive Permissions Revoked	Average Time to Rotate NHI Credentials
NHI-Related Incident Rate Decreased	Percentage of NHIs without Assigned Ownership Detected

> **Non-human identity is a top priority for many teams and it is challenging to gain insight into the full scale and scope of the issue. Token provides an easy solution to get that visibility to then determine where you need to take action.**
>
> **Mandy Andress,** CISO at Elastic

token

08

HOW TOKEN SECURITY IS SOLVING THE AI AGENT AND NHI SECURITY CRISIS

Securing your cloud and AI environments begins with protecting both human and non-human identities that access your critical systems. Token Security has the answer to solving this AI agent and identity crisis with a machine-first identity security approach.

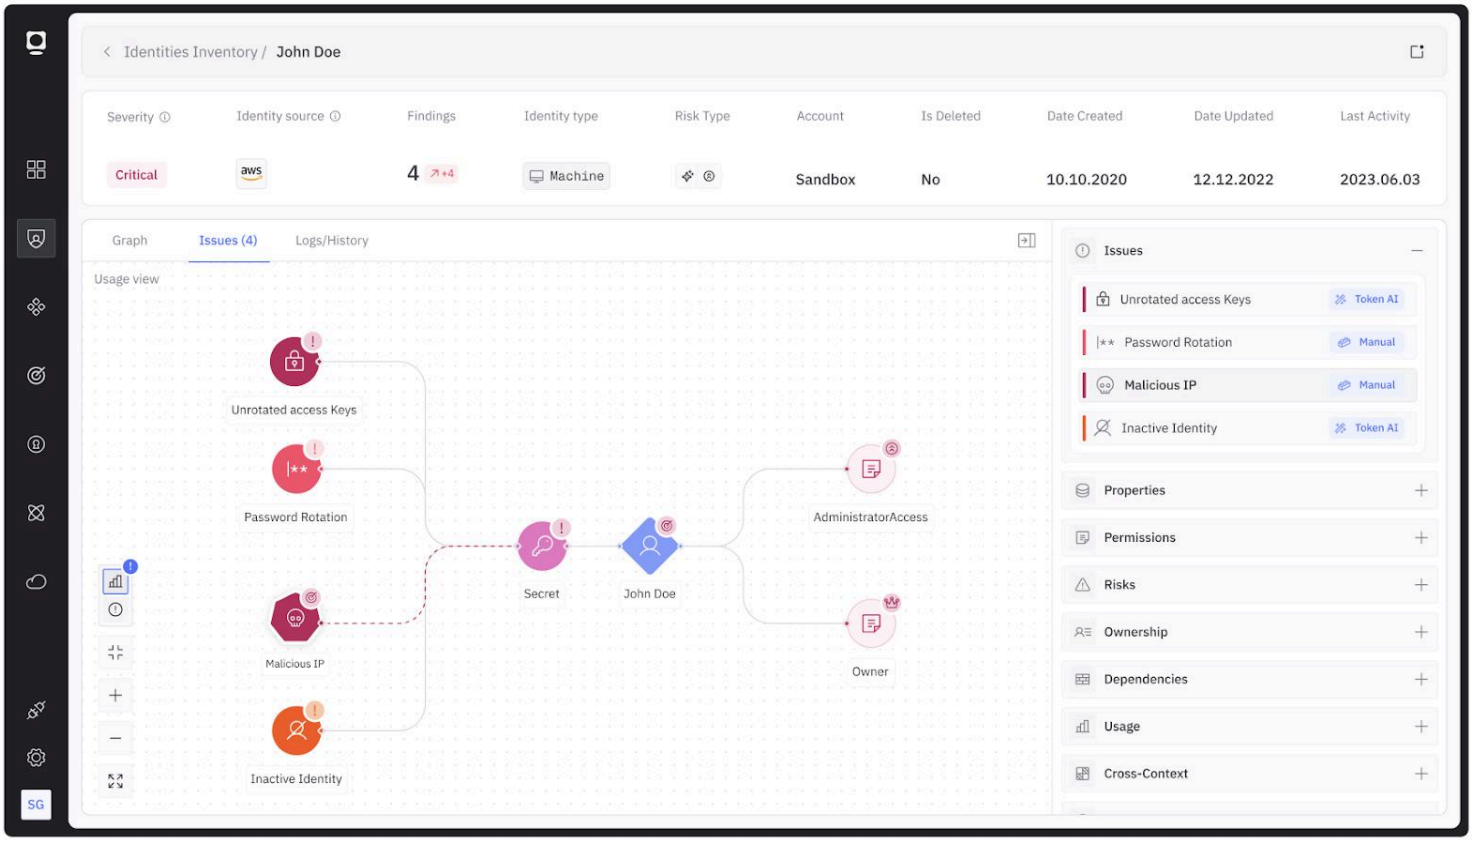

Discover all AI agents and identities mapped across your cloud environments with ownership attribution behind each identity. Token Security provides AI agent and machine identity lifecycle control to successfully deactivate offboarded employee accounts, rotate keys and credentials, and remove stale agents and identities. Manage the entire lifecycle at every stage of the journey.

Prioritize risk remediation efforts based on contextual business information, so your team knows where to take action first and do it with confidence. Focus on what matters most and demonstrate clear ROI by preventing threats.

Discovery	Visibility	Prioritize	Remediate
Reveal all AI agents and non-human identities in your hybrid, multi-cloud environment	Gain full visibility over agent and identity ownership and usage	Prioritize risks and remediate exposures with contextualized information	Remediate critical risks with intelligence and confidence

> **Token Security has enhanced our security by providing us accurate and relevant visibility into machine identity risks, something I have yet to see from other vendors. It's the first service I've seen that delivers on the widely-made claim of finding machine identity risk. It provides us with the right amount of detail to mitigate issues without burdening us with tens of rows of useless alerts.**
>
> **Jonathan Jaffe,** CISO at Lemonade

Discover how Token Security's machine-first approach helps solve your Agentic AI and NHI security challenges
Let's Talk