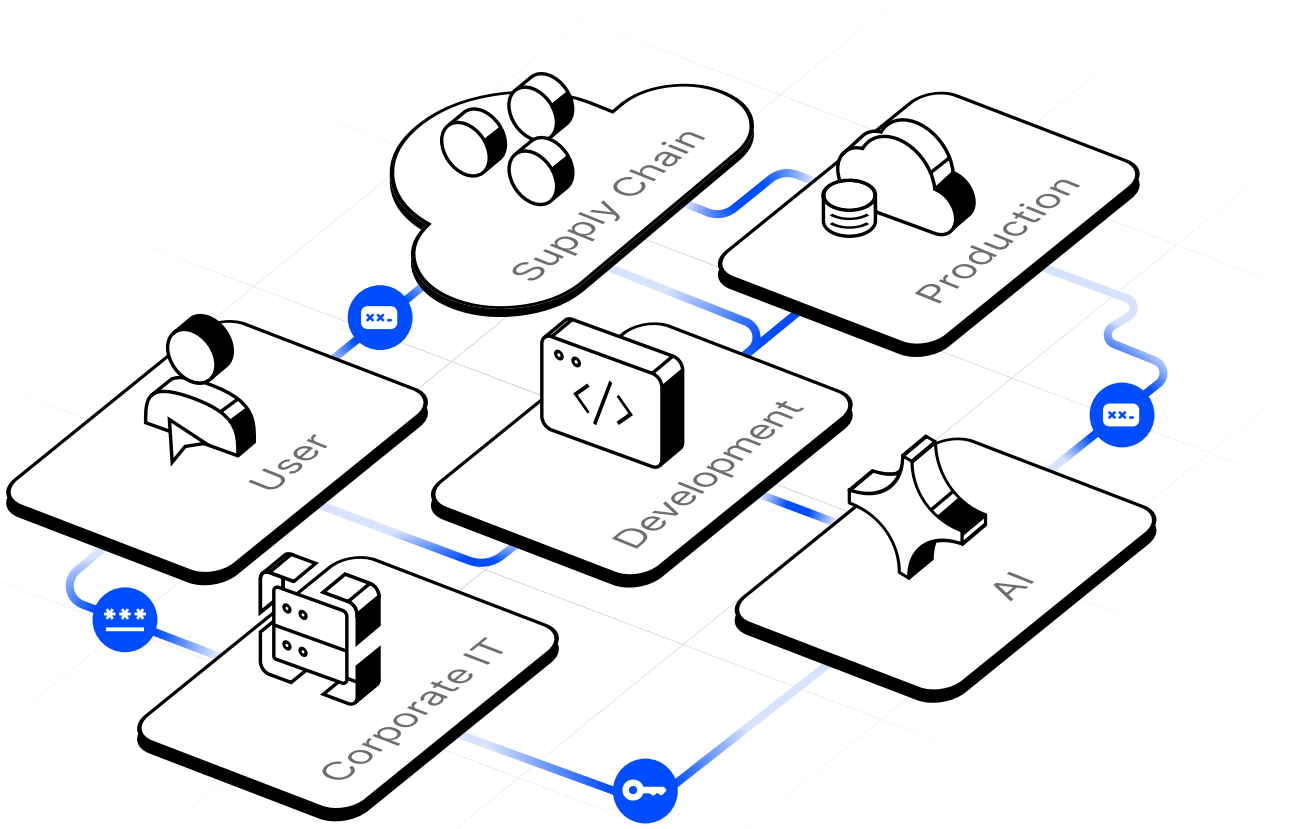
clutch



Mapping the Enterprise Non-Human Identity Attack Surface

A Strategic Security Analysis

Table of Contents

Executive Summary 3

The Scope of the NHI Security Challange 4

- User Domain: The Human-Machine Interface 5
- Corporate IT Domain: The Operational Backbone 6
- Supply Chain Domain: The Extended Enterprise 7
- Development Domain: The Innovation Engine 8
- Production Domain: The Mission-Critical Environment 9
- AI Domain: The Emerging Intelligence Frontier 10

Domain Risk Assessment: Strategic Security Priorities 11

- Corporate IT Domain 12
- Production Domain 13
- User Domain 14
- Supply Chain Domain 15
- Development Domain 16
- AI Domain 17

Required NHI Security Capabilities 19

- Discovery and Inventory 19
- Lifecycle Management and Governance 19
- Posture Management 20
- Threat Detection 21
- Incident Response and Recovery 21

Strategic Implementation Framework 22

- Immediate Actions (0-90 Days): Establishing Visibility 22
- Short-Term Initiatives (3-12 Months): Building Governance 22
- Long-Term Strategy (12+ Months): Achieving Maturity 23

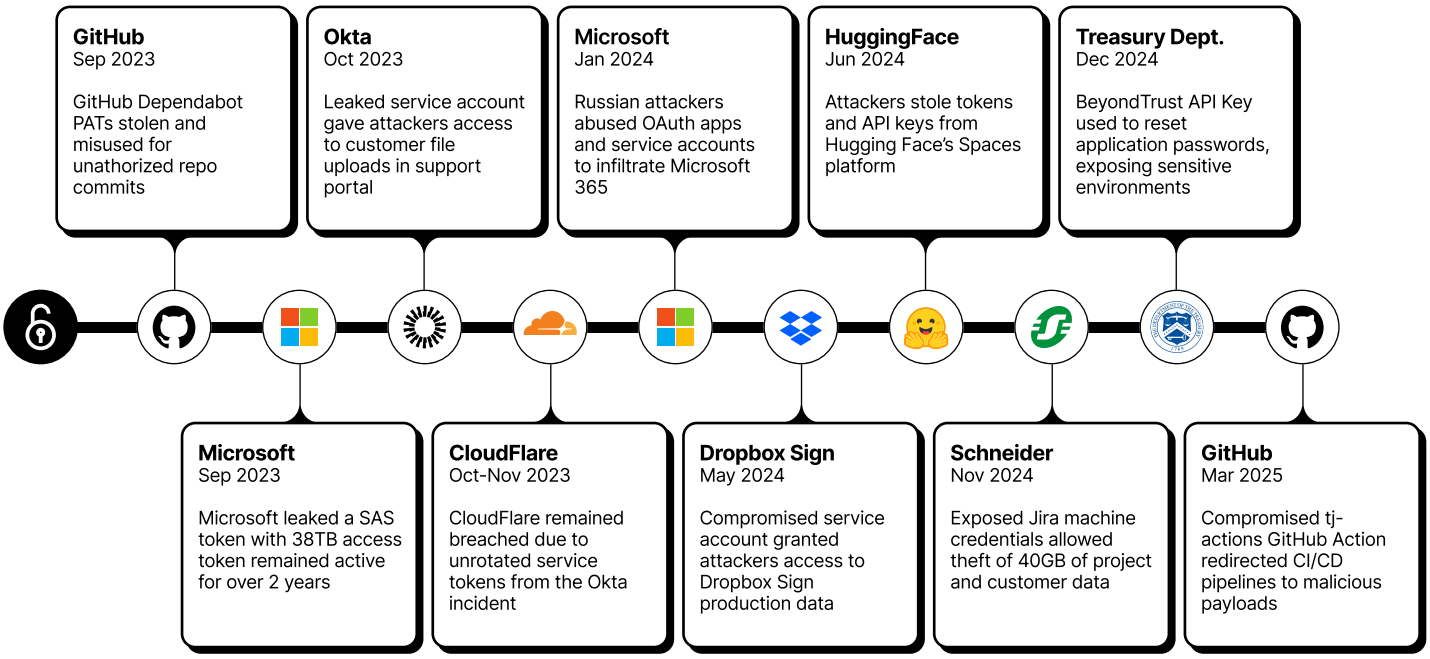
Conclusion: The Imperative for Action 24

Executive Summary

Machine identities—colloquially referred to in the industry as Non-Human Identities (NHIs), including API keys, secrets, tokens, service accounts, and certificates—outnumber human identities 82:1 in modern enterprises yet receive substantially fewer security controls. This disparity has created an expanding attack surface that threat actors are increasingly exploiting.

Our analysis identifies six enterprise domains where NHIs operate: User, Corporate IT, Supply Chain, Development, Production, and AI. The AI and Development domains present the highest risk due to immature security practices and rapid credential proliferation, while Corporate IT represents the most mature and manageable domain.

Recent high-profile incidents demonstrate that NHI compromise is a current reality, not a theoretical risk. While organizations invest heavily in human identity security with MFA and behavioral monitoring, machine identities typically operate with static credentials, minimal governance, and limited oversight.



Notable breaches involving Non-Human Identities

Key Recommendations

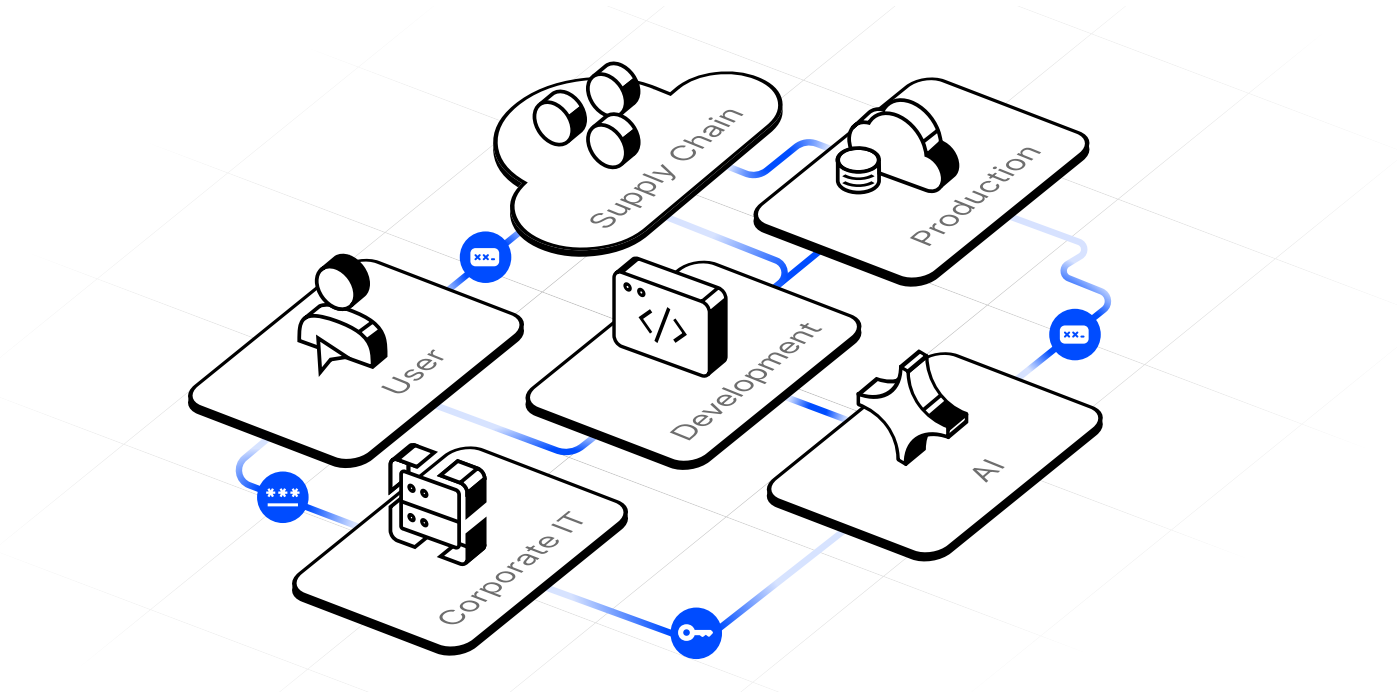
- Implement comprehensive NHI discovery across all enterprise domains within 90 days
- Prioritize AI and Development domain security initiatives based on risk assessment
- Initiate lifecycle management and behavioral monitoring for critical NHIs
- Establish automated incident response capabilities for credential compromise

Organizations that begin systematic NHI security initiatives will be significantly better positioned to defend against this evolving threat landscape while maintaining operational efficiency.

The Scope of the NHI Security Challenge

To understand the magnitude of the NHI security crisis, we must first map where these identities exist and operate within modern enterprises. The enterprise environment can be conceptualized as six distinct operational domains, each serving specific business functions and different stakeholders while harboring unique populations of Non-Human Identities.

These domains - User Domain, Corporate IT Domain, Supply Chain Domain, Development Domain, Production Domain, and AI Domain - all represent the complete digital ecosystem where NHIs proliferate. Each domain operates with different risk profiles, stakeholder dynamics, and security maturity levels, creating a complex attack surface that requires domain-specific understanding and security approaches.



Understanding these domains is critical, because the security pendulum has swung dramatically: attackers have shifted their focus from targeting well-defended human identities to exploiting the vast, under-secured populations of machine identities that enable modern digital operations. This shift is particularly dangerous because NHIs are typically privileged by design - a single compromise can provide Lateral Movement and Privilege Escalation opportunities with a potentially devastating impact.

Moreover, this trend will only worsen. Security teams lack comprehensive visibility across these diverse domains, and rely on tools that were never designed to secure machine identities. Meanwhile, engineering teams, end users, and IT departments continue to embrace new technologies and tools in pursuit of business efficiency and velocity. This creates a dangerous paradox: while business velocity accelerates and the enterprise risk profile grows exponentially, security teams are left blind - lacking appropriate tooling, visibility, and expertise to govern the expanding NHI population they're meant to protect.

The following pages break down each domain, highlighting NHI types, key actors, and security risks.

User Domain

The Human-Machine Interface

Domain Overview

The User Domain represents the digital workspace where employees interact with enterprise systems to perform their daily functions. This domain encompasses web browsing, email communication, SaaS application usage, file sharing, and collaboration platforms.

Main Actors

End-users who, through their productivity activities, generate and interact with numerous Non-Human Identities; IT and Security teams managing and securing user access and monitoring user behavior.

NHI Population

- OAuth applications authorized by users for productivity and integration
- Browser-stored API keys and session tokens
- SaaS application credentials managed by individual users

Characteristics

The User Domain presents a particularly challenging attack surface because security controls must balance usability with protection. End-users regularly install applications, grant permissions, and store credentials without full understanding of the security implications. The domain's distributed nature—spanning corporate devices, BYOD endpoints, and cloud services—creates multiple attack vectors.

Key Vulnerabilities

OAuth Consent Abuse: Users can be socially engineered into granting excessive permissions to malicious applications, creating persistent access tokens with broad enterprise access.

Credential Harvesting: Malware on endpoints can extract stored tokens and API keys from browsers and applications.

Shadow IT Proliferation: Unmanaged application installations create ungoverned access pathways that bypass corporate security controls.

Token Persistence: OAuth tokens and API keys often remain active indefinitely, providing long-term access even after business needs change.

Corporate IT Domain

The Operational Backbone

Domain Overview

The Corporate IT Domain encompasses the traditional enterprise infrastructure that supports business operations. This domain includes on-premises servers, network infrastructure, directory services, legacy applications, and hybrid cloud connections.

Main Actors

IT administrators, system engineers, network operations teams, and automated processes that maintain enterprise services; security teams responsible for infrastructure protection.

NHI Population

- Active Directory service accounts for system automation
- Network device credentials (switches, routers, firewalls)
- Legacy application service accounts with embedded credentials
- SSH keys for server access and automation
- PKI certificates for system-to-system authentication

Characteristics

The Corporate IT Domain often represents the oldest and most established part of the enterprise attack surface. Legacy systems in this domain frequently lack integration with modern identity management platforms, relying instead on static credentials and shared accounts for operational continuity.

Key Vulnerabilities

- **Credential Reuse:** Static passwords and keys used across multiple systems create lateral movement opportunities.
- **Privilege Accumulation:** Service accounts collect permissions over time without regular review or cleanup.
- **Legacy System Gaps:** Older systems bypass modern security controls, operating with minimal logging and monitoring.
- **Hardcoded Credentials:** Cleartext credentials embedded in scripts and configuration files stored on file shares and legacy systems.

Supply Chain Domain

The Extended Enterprise

Domain Overview

The Supply Chain domain represents the complex web of vendor relationships, SaaS integrations, and partner connections that extend enterprise capabilities. This domain includes cloud service providers, SaaS applications, managed service providers, and business partners.

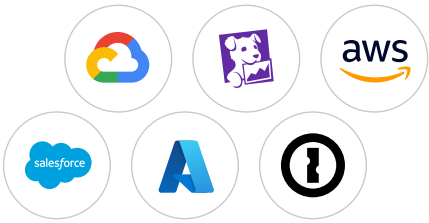
Main Actors

External vendors, integration platforms, automated systems that facilitate business relationships; vendor management teams; and security teams responsible for third-party risk assessment.



NHI Population

- OAuth tokens granted to SaaS vendors for integration purposes
- API keys shared with third-party service providers
- Partner service accounts for B2B integrations
- Cloud service provider credentials for managed services
- Vendor-managed certificates and encryption keys



Characteristics

The Supply Chain domain presents unique challenges because organizations must balance operational necessity with security control. Enterprises often lack visibility into how external parties manage shared credentials, creating blind spots in their security posture. This domain has become a primary target for threat actors, with numerous supply chain-driven NHI attacks demonstrating how vendor compromises can cascade across entire customer ecosystems.

Key Vulnerabilities

- **Vendor Credential Mismanagement:** Third parties may fail to implement proper credential rotation, storage, or access controls.
- **Integration Sprawl:** Complex chains of integrations create cascading access relationships that are difficult to map and secure.
- **Trust Relationship Persistence:** Vendor access often continues beyond the business relationship, creating dormant but active attack vectors.
- **Limited Revocation Capability:** Organizations may lack the ability to quickly revoke third-party access across all integrated systems.

Development Domain

The Innovation Engine

Domain Overview

The Development Domain encompasses the software development lifecycle, from code creation through deployment. This domain includes source code repositories, CI/CD pipelines, development environments, and deployment automation.

Main Actors

Software developers, DevOps engineers, automated build and deployment systems; platform engineering teams; and security teams responsible for application security and DevSecOps.

NHI Population

- Personal Access Tokens for repository access and automation
- Secrets stored in pipelines or CI/CD plugins
- Container registry credentials for image storage and retrieval
- Cloud deployment credentials embedded in automation scripts
- API keys for third-party development services and tools
- Secrets embedded in application code and configuration files

Characteristics

The Development Domain operates under high velocity and innovation pressure, often prioritizing speed and convenience over security. The domain's distributed nature - spanning developer workstations, cloud repositories, and deployment pipelines - creates numerous opportunities for credential exposure.

Key Vulnerabilities

- **Secret Sprawl:** Credentials scattered across repositories, configuration files, and deployment scripts, often in plaintext.
- **Version Control Exposure:** Git history preserves deleted secrets indefinitely, creating long-term exposure risks.
- **Pipeline Privilege Escalation:** CI/CD accounts often hold broad permissions across multiple environments for deployment automation.
- **Developer Convenience Practices:** Hardcoding credentials and sharing keys for development efficiency creates persistent security gaps.

Production Domain

The Mission-Critical Environment

Domain Overview

The Production Domain represents the live operational environment where business-critical applications and services run. This domain includes cloud infrastructure, containerized workloads, microservices architectures, and production databases.

Main Actors

Production systems, automated orchestration platforms, operational monitoring tools; Site Reliability Engineers (SREs); and security teams responsible for production security monitoring.

NHI Population

- Cloud IAM roles and access keys for infrastructure automation
- Kubernetes service accounts for container orchestration
- Workload identity federation tokens for cross-cloud authentication
- Instance metadata credentials for cloud resource access
- Microservices authentication tokens for service-to-service communication
- Database connection credentials and application service accounts

Characteristics

The Production Domain demands the highest levels of availability and performance, often creating tension between security controls and operational requirements. The domain's dynamic nature - with auto-scaling, ephemeral workloads, and constant deployment - generates and destroys credentials automatically. While modern cloud platforms offer ephemeral credentials like AWS IAM roles, Azure Managed Identities, and GCP Workload Identities, many organizations still rely on legacy static credentials like access keys and client secrets that are long-lived and harder to manage. Additionally, most security tools focus on the infrastructure layer rather than the identity layer, creating visibility gaps in credential usage and lifecycle management.

Key Vulnerabilities

- **Over-Privileged Service Accounts:** Default cloud configurations often grant excessive permissions for operational convenience.
- **Metadata Service Exploitation:** Compromised workloads can access cloud credentials through instance metadata services.
- **Lateral Movement Opportunities:** Service mesh architectures enable authenticated traversal between compromised services.
- **Static Credential Persistence:** Long-lived access keys and client secrets that lack proper rotation and lifecycle management.

AI Domain

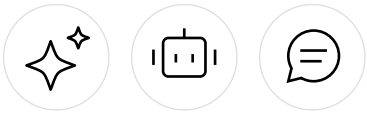
The Emerging Intelligence Frontier

Domain Overview

The AI Domain represents the newest and rapidly expanding domain where artificial intelligence and automation systems operate. This domain includes Large Language Models (LLMs), AI agents, automation platforms, and machine learning pipelines. The most significant development in this domain is the rapid adoption of AI agents by enterprises, often with limited understanding of the inherent security risks.

Main Actors

AI agents that interact with enterprise data and systems; AI/ML engineers; and data scientists.



NHI Population

- API keys for AI service providers (OpenAI, Anthropic, etc.)
- Service accounts for AI training data access
- NHIs used by AI agents for autonomous system interactions
- Machine learning pipeline credentials for data processing
- Hardcoded secrets inadvertently included in training datasets



Characteristics

The AI Domain presents unique security challenges because AI systems require broad access to organizational data for training and operation, while automation platforms accumulate extensive permissions for workflow execution. This domain is characterized by rapid innovation and relatively immature security practices. The most critical concern is AI agent proliferation - autonomous systems that use NHIs to perform operations across enterprise systems, many with write privileges, creating unprecedented attack surfaces.

Key Vulnerabilities

- Data Exposure Through Training: LLMs may inadvertently learn and reproduce credentials from training datasets.
- AI Agent Manipulation: Social engineering attacks against AI-powered helpdesk or automation systems to perform unauthorized actions.
- Automation Privilege Accumulation: AI-driven automation platforms often acquire broad permissions across multiple enterprise systems.
- Agent Sprawl and NHI Proliferation: Rapid deployment of AI agents creates sprawling populations of machine identities, each requiring authentication credentials to perform autonomous operations, often with elevated privileges that enable cross-system access and modification capabilities.

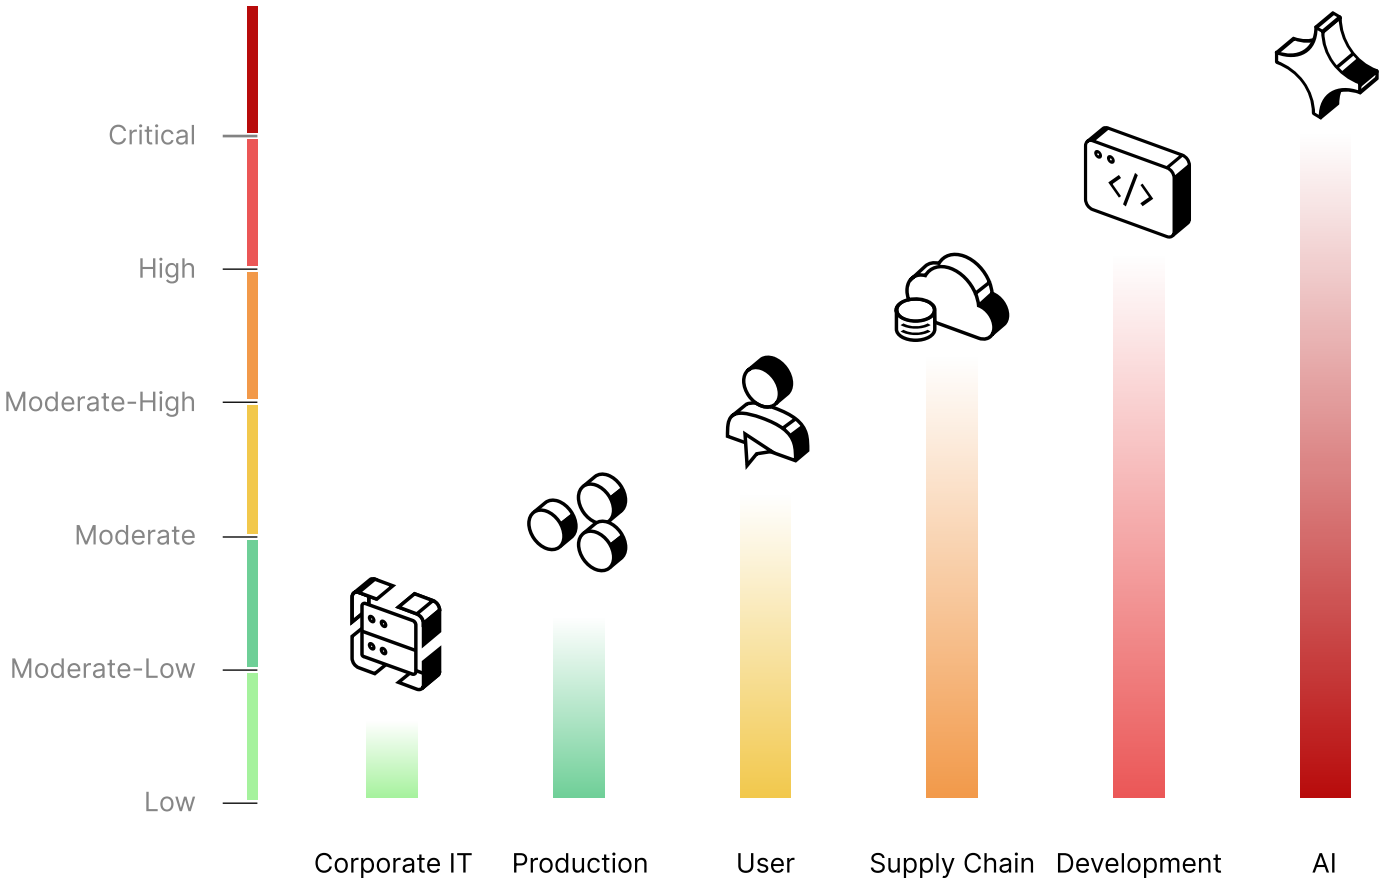
Domain Risk Assessment: Strategic Security Priorities

Understanding the relative risk levels across enterprise domains is critical for prioritizing security investments and resource allocation. While all six domains present significant NHI security challenges, they vary substantially in their risk profiles based on several key factors that determine both the likelihood of successful attacks and the potential business impact.

To provide actionable guidance for security leaders, we have assessed each domain against four critical risk parameters:

1. **Security Tooling Maturity:** Availability of purpose-built security solutions
2. **Governance Complexity:** Difficulty of implementing and enforcing consistent policies
3. **Attack Surface Size:** Volume and diversity of NHI populations
4. **Blast Radius Potential:** Scope of access and impact if compromise occurs

These parameters collectively determine both the probability of successful exploitation and the magnitude of potential damage.



The following analysis ranks enterprise domains from lowest to highest risk, providing the strategic context needed for risk-based security planning:

Corporate IT Domain

Risk Assessment

Low

Rationale

The Corporate IT Domain represents the most mature and controllable domain from a security perspective, benefiting from decades of enterprise security investment and established operational practices.

Factor	Type	Decription
Legacy System Integration	Risk	Older systems may lack modern security control integration
Security Tooling Maturity	Mitigation	Decades of investment in Active Directory, PAM solutions, network monitoring, and identity governance platforms
Governance Complexity	Mitigation	Well-understood RBAC models, established approval workflows, and mature audit processes
Attack Surface Size	Mitigation	Relatively contained with clearly defined infrastructure boundaries
Blast Radius Potential	Mitigation	The network perimeter, additional segmentation, and role-based access controls contain compromise scope
Regulatory Compliance	Mitigation	Compliance requirements enforce baseline security hygiene



Production Domain

Risk Assessment

Moderate Low

Rationale

The Production Domain benefits from modern cloud-native security capabilities and high business visibility, though it requires careful management of privileged access requirements.

Factor	Type	Decription
Attack Surface Size	Risk	Substantial due to microservices architectures and dynamic scaling
Static Credential Usage	Risk	Many organizations still use long-lived access keys and client secrets
Security Tooling Maturity	Mitigation	Modern cloud IAM systems with built-in least privilege capabilities
Governance Complexity	Mitigation	Infrastructure-as-Code enables consistent security control deployment
Blast Radius Potential	Mitigation	Granular cloud IAM and microservices boundaries limit lateral movement
Business Visibility	Mitigation	Production issues receive immediate attention and resources



User Domain

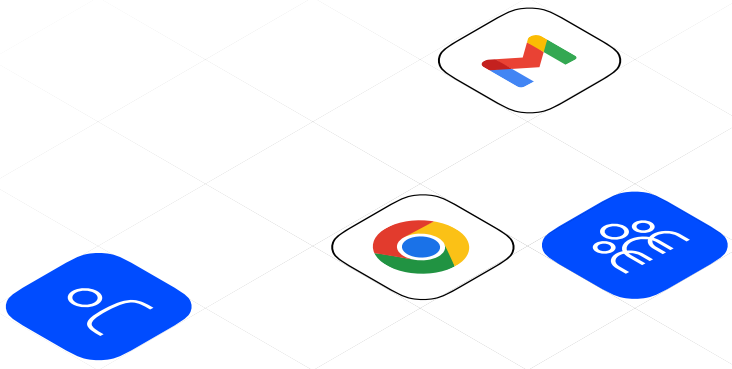
Risk Assessment

Moderate

Rationale

The User Domain presents moderate risk due to its distributed nature and social engineering susceptibility, but benefits from naturally limited access scope and improving security tooling.

Factor	Type	Decription
Attack Surface Size	Risk	Every user generates multiple NHIs through daily activities
Social Engineering Susceptibility	Risk	Users can be manipulated into granting excessive permissions
Security Tooling Maturity	Mitigation	CASB, EDR, Secure Browsers, and ZTNA tools providing increasing visibility and control
Blast Radius Potential	Mitigation	User-generated NHIs inherit limited creator permissions
Security Awareness	Mitigation	Training programs address social engineering risks



Supply Chain Domain

Risk Assessment

Moderate-High

Rationale

The Supply Chain Domain presents elevated risk due to limited organizational control over external security practices and complex trust relationships that can cascade compromise impact.

Factor	Type	Decription
Security Tooling Maturity	Risk	Limited control over third-party credential management practices
Governance Complexity	Risk	Challenge of enforcing policies across diverse vendor relationships
Attack Surface Size	Risk	Growing reliance on SaaS integrations and cloud service providers
Blast Radius Potential	Risk	Highly variable - some vendors have broad enterprise access
Trust Relationship Persistence	Risk	Vendor access often continues beyond business relationships
Vendor Risk Programs	Mitigation	API gateways and contract security requirements provide partial control



Development Domain

Risk Assessment

High

Rationale

The Development Domain represents high risk due to the fundamental tension between development velocity and security controls, combined with the distributed nature of modern software development practices.

Factor	Type	Decription
Security Tooling Maturity	Risk	Secret scanning tools remain fragmented across development lifecycle
Governance Complexity	Risk	Developer autonomy requirements conflict with security controls
Attack Surface Size	Risk	Massive distribution of secrets across repositories and pipelines
Blast Radius Potential	Risk	Development credentials often span multiple environments
Version Control Persistence	Risk	Git history preserves leaked secrets indefinitely
Developer Practices	Risk	Convenience often prioritized over security, with hardcoded secrets as a direct result



AI Domain

Risk Assessment

Critical

Rationale

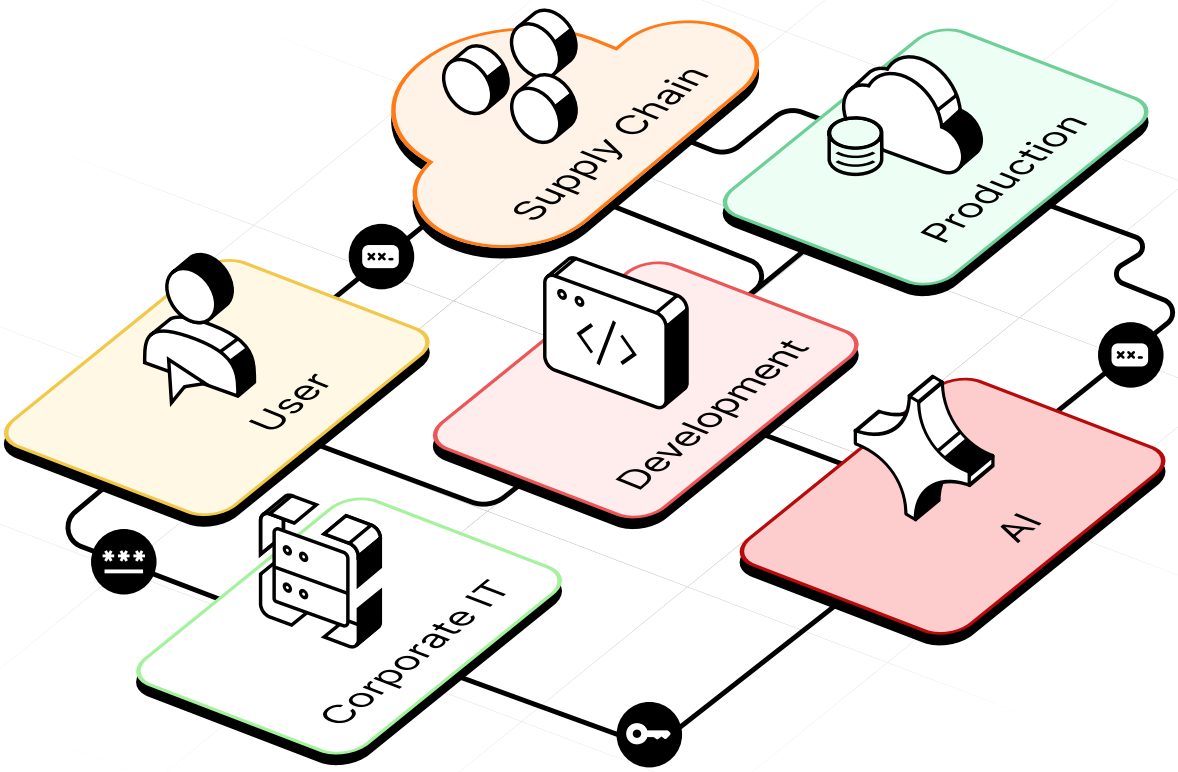
The AI domain represents the highest-risk area due to nascent security practices, broad access requirements, and the pace of enterprise adoption outpacing governance. Its inherent unknown risks and lack of mature controls make it especially vulnerable in existing security frameworks.

Factor	Type	Decription
Security Tooling Maturity	Risk	Minimal enterprise security solutions for AI-specific risks
Governance Complexity	Risk	Few organizations have AI-specific security policies
Attack Surface Size	Risk	Rapidly expanding AI deployment without security assessment
Blast Radius Potential	Risk	AI systems granted broad organizational data access
Credential Learning	Risk	LLMs may learn and reproduce credentials from training data
Agent Proliferation	Risk	AI agents create sprawling NHI populations with elevated privileges

Strategic Implication

The AI Domain represents an *"attack surface explosion"* occurring faster than security practices can adapt, requiring immediate attention and investment to prevent widespread organizational exposure.





Low Risk Critical Risk

This risk assessment provides the foundation for strategic security planning, enabling organizations to prioritize their NHI security investments based on both current threat levels and emerging risk trajectories. Organizations should focus immediate attention on AI Domain governance and Development Domain secret management while leveraging existing strengths in Corporate IT Domain to establish security frameworks for newer domains.

Required NHI Security Capabilities

Based on the domain analysis and risk assessment, organizations require a comprehensive security framework that addresses the unique challenges of non-human identity management. This framework must span five core capability areas, each addressing specific aspects of the NHI security lifecycle.



Discovery and Inventory

The Foundational Requirement

Organizations cannot secure what they cannot see. Comprehensive NHI discovery forms the foundation of any effective security program, providing the visibility necessary for all subsequent security activities.

Core Capabilities

- **Multi-Domain Discovery:** Real-time scanning across cloud environments, SaaS platforms, on-premises infrastructure, code repositories, and emerging AI platforms, with ongoing discovery processes that detect new NHIs as they are created and identify changes to existing identities.
- **Comprehensive Inventory Management:** Centralized repository of all discovered NHIs with metadata including creation date, last usage, permissions, and business context.
- **Identity Correlation:** Automated mapping of the same identity across different systems and contexts to understand the complete access footprint.
- **Business Context Attribution:** Linking each NHI to its business purpose, ownership, and operational requirements to enable risk-based decision making.

Implementation Approach

Deploy API-based discovery connectors across all enterprise domains with automated classification engines that can distinguish between different NHI types and their associated risk profiles.



Lifecycle Management and Governance

The Core Control Requirement

NHIs require structured governance processes that manage their entire lifecycle from creation through decommissioning, ensuring that access remains aligned with business requirements.

Core Capabilities

- **Ownership Attribution:** Clear assignment of responsibility for each NHI to specific individuals or teams who can make access decisions.
- **Attestation and Recertification:** Regular review processes where owners validate the continued need for each NHI and its current permissions.
- **Approval Workflows:** Structured processes for NHI creation, modification, and renewal that include appropriate business and security reviews.
- **Automated Lifecycle Policies:** Rules-based systems that enforce expiration dates, rotation schedules, and decommissioning procedures.

Implementation Approach

Establish policy-driven governance engines integrated with existing IT service management platforms, with automated enforcement capabilities and exception handling processes.



Posture Management

The Optimization Requirement

Organizations must continuously assess and improve their NHI security posture, identifying overprivileged accounts, risky configurations, and opportunities for access reduction.

Core Capabilities

- **Privilege Analysis:** Continuous evaluation of NHI permissions against usage patterns to identify over-privileged accounts.
- **Risk Scoring:** Dynamic assessment of each NHI's risk based on permissions, usage patterns, access scope, and business context.
- **Configuration Monitoring:** Detection of security configuration drift and non-compliance with established security policies.
- **Proactive Least Privilege Enforcement:** Automated reduction of permissions to minimum required levels and restriction of access scope based on actual usage patterns.
- **Automated Remediation:** Capability to automatically implement security improvements such as permission reduction and access restriction.

Implementation Approach

Deploy analytics platforms that can baseline normal NHI behavior and identify deviations, with integration to policy enforcement systems for automated remediation.



Threat Detection

The Protection Requirement

Organizations need real-time visibility into NHI usage patterns to detect unauthorized access, credential theft, and other malicious activities.

Core Capabilities

- **Behavioral Baseline Establishment:** Learning normal usage patterns for each NHI to enable accurate anomaly detection.
- **Real-Time Monitoring:** Continuous observation of NHI activities across all enterprise domains with intelligent alerting.
- **Cross-Domain Correlation:** Ability to connect related activities across different enterprise domains to identify complex attack patterns.
- **Integration with Security Operations:** Seamless connection with SIEM, SOAR, and incident response platforms for coordinated threat response.

Implementation Approach

Implement machine learning-based behavioral analysis platforms with customizable detection rules and automated integration with existing security operations workflows.



Incident Response and Recovery

The Resilience Requirement

When NHI compromises occur, organizations need rapid response capabilities that can contain the threat, assess the impact, and restore secure operations.

Core Capabilities

- **Blast Radius Analysis:** Immediate assessment of which systems and data may be affected by a compromised NHI.
- **Automated Containment:** Rapid credential rotation, access revocation, and system isolation capabilities.
- **Forensic Data Collection:** Comprehensive logging and audit trail capabilities that support incident investigation and legal requirements.
- **Recovery Orchestration:** Automated processes for restoring secure operations with new credentials and validated system integrity.

Implementation Approach

Implement machine learning-based behavioral analysis platforms with customizable detection rules and automated integration with existing security operations workflows.

Strategic Implementation Framework

Immediate Actions (0-90 Days)

Establishing Visibility

The first priority must be gaining comprehensive visibility into the existing NHI population. Organizations should:

1. **Conduct Multi-Domain Discovery:** Deploy discovery tools across all six enterprise domains to identify and catalog existing NHIs.
2. **Eliminate Critical Exposures:** Remove hardcoded secrets from code repositories and rotate credentials with known exposure risks.
3. **Establish Basic Monitoring:** Implement logging and alerting for critical service accounts and high-privilege OAuth applications.

Short-Term Initiatives (3-12 Months)

Building Governance

With visibility established, organizations can implement structured governance and risk management:

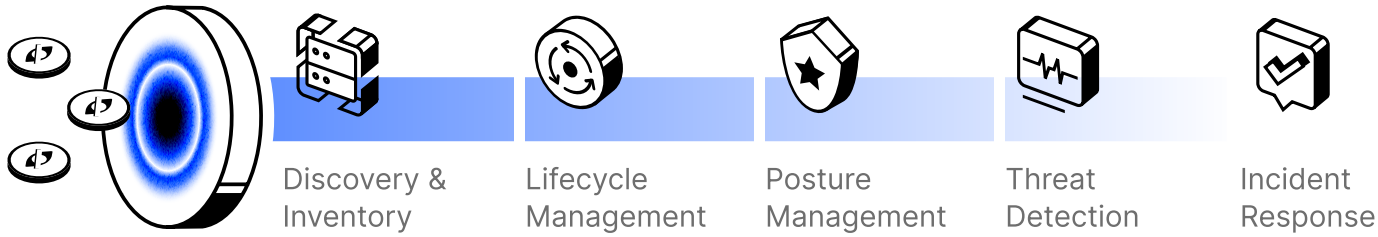
1. **Deploy Governance Framework:** Establish ownership attribution, approval workflows, and lifecycle management policies.
2. **Implement Risk-Based Controls:** Deploy privilege analysis and automated least-privilege enforcement for high-risk NHIs.
3. **Enhance Detection Capabilities:** Implement behavioral monitoring and anomaly detection across all NHI populations.

Long-Term Strategy (12+ Months)

Achieving Maturity

The final phase focuses on automation, optimization, and advanced threat response:

1. **Automated Security Orchestration:** Deploy AI-driven security automation for autonomous threat detection and response.
2. **Zero Trust Architecture:** Implement continuous verification and adaptive access controls for all NHIs.
3. **Supply Chain Integration:** Establish comprehensive third-party NHI risk management and monitoring programs.



Conclusion: The Imperative for Action

The evidence is clear: Machine Identities represent the largest and fastest-growing attack surface in modern enterprises. The analysis presented in this paper demonstrates that NHI compromise is not a theoretical future threat—it is happening now, with severe business consequences for organizations across every industry.

The fundamental asymmetry between human and non-human identity security creates an environment where attackers can operate with minimal detection risk while maintaining persistent access to critical enterprise resources. This situation is unsustainable in an era where digital transformation has made NHIs essential to business operations.

Organizations that continue to treat NHIs as operational afterthoughts rather than critical security assets are providing attackers with authenticated pathways to their most valuable data and systems. The cost of reactive security - measured in breach response, regulatory penalties, business disruption, and competitive disadvantage - far exceeds the investment required to implement proactive NHI security.

The security capabilities and implementation framework outlined in this paper provide a proven path forward, but the window for proactive action is narrowing as attackers increasingly focus on these under-defended assets. Organizations must act decisively to secure their Non-Human Identities before they become the next headline breach.

About Clutch Security

Clutch Security specializes in securing Non-Human Identities across enterprise environments. As Agentic AI fundamentally changes how NHIs operate within organizations, Clutch provides the visibility, governance, and risk management capabilities required for safe AI adoption at scale.

