



Non-Human Identity
Management Group

NHI MGMT GROUP • LIFECYCLE

NHI & AI Agent Lifecycle Management

Without lifecycle management, non-human identities and AI agents become security blind spots. Every stage matters — from provisioning and discovery to securing credentials, monitoring and prevent controls.

AI agents are non-human identities that decide and act on their own, at machine speed. The same seven stages apply, with additional controls at each one.

Each stage covers non-human identities and AI agents — see page 2.





What each stage covers — for NHIs and AI agents

STAGE	NON-HUMAN IDENTITIES	AI AGENTS
1 Provisioning & Decommissioning	• Deliver static NHIs straight into a secrets vault • Prefer just-in-time NHIs with least privilege • Recertify regularly; decommission and rotate on leavers	• Register each agent at creation, with an owner and purpose • One identity per agent — never borrowed credentials • Revoke agent tokens and keys when the task ends
2 Discovery & Inventory	• Inventory IdPs, directories and local accounts everywhere • Scan code, wikis and file shares for plain-text secrets • Map each found secret to its account and usage	• Discover agents, MCP servers and AI API keys • Scan prompts and agent configs for embedded secrets • Keep an agent registry: owner, purpose, access
3 Classification	• Assign owners to applications, not people • Record privilege level and breadth of access • Flag inactive accounts and hard-coded NHIs	• Record autonomy: suggest, approved or fully autonomous • Map delegated access: who the agent acts for • Flag agents that can change data or touch production
4 Posture Management	• Cut excessive permissions to least privilege • Remove inactive accounts and fix misconfigurations • No shared NHIs; separate prod and non-prod	• Review agent tool permissions and OAuth scopes • Find long-lived or over-scoped agent tokens • One identity per agent, per environment
5 Securing Credentials	• Move plain-text secrets into a vault or encryption • Use passwordless options such as certificates • Automate rotation once dependencies are mapped	• Issue short-lived, task-scoped agent credentials • Keep secrets out of prompts, logs and agent memory • Rotate AI API keys; revoke on any exposure
6 Monitoring Controls	• Use behaviour analytics to spot anomalies at scale • Detect humans using NHIs • Tune detections to keep false positives low	• Log every agent action and tool call • Alert on unusual agent behaviour • Watch for prompt injection misusing agent access
7 Prevent Controls	• Block code check-ins that contain secrets • Move to dynamic, ephemeral secrets • Enforce policy in real time, e.g. identity control plane	• Enforce real-time guardrails on agent actions • Keep a tested kill switch for every agent • Require human approval for high-risk actions

Where are you today? Take the free NHI & AI Governance Maturity Assessment at nhimg.org/nhi-ai-maturity-assessment · Full guide: nhimg.org/nhi-lifecycle-management-guide