KEYFORGEAI

WHITEPAPER

Governing *Non-Human Identity* at Scale

Service accounts, secrets, API keys, and machine identities now outnumber humans. A practical blueprint for bringing them under continuous governance.

JUNE 2026 · WWW.KEYFORGE.AI

EXPLAINABLE · AUDITABLE · POLICY-DRIVEN

The Continuous Identity Reasoning Platform

Ask an identity team how many identities they govern and they will quote their headcount. Ask how many service accounts, API keys, secrets, and machine identities exist across the same estate, and the honest answer is almost always: nobody knows. The ratio has quietly inverted — and the tooling has not followed.

§ 01 — THE PROBLEM

The population that nobody owns

Non-human identities accumulate. Unlike employees, they do not attend training, do not respond to certification emails, and do not leave when their project ends. A service account created for a 2019 integration is very likely still active, still privileged, and still using the credential it was born with.

The governance consequences are concrete: orphaned accounts with no owner to ask; shared credentials passed between teams; API keys with scopes granted for a use case that no longer exists; secrets embedded in code and pipelines, invisible to any review process. Every one of these is an audit finding waiting to be written and an attack path waiting to be used.

§ 02 — WHY IGA MISSES IT

Built for the employee lifecycle

Traditional IGA is anchored to the HR record: joiner, mover, leaver. Non-human identities have no HR record, no manager, and no termination date. They enter the estate through infrastructure automation, application onboarding, and developer tooling — paths that never touch the request catalog.

Quarterly certification makes the mismatch worse. A campaign that reviews a service account once a year is reviewing a snapshot; the account's usage, scope, and blast radius change weekly. Event-driven governance cannot govern a population whose events happen outside the governance system.

> *An identity program that governs only humans governs a shrinking minority of its own identities — and a shrinking minority of its own risk.*

§ 03 — THE BLUEPRINT

Five steps to continuous NHI governance

Step 1 — Inventory relentlessly. Aggregate from directories, cloud IAM, vaults, CI/CD systems, and application service-account tables. Expect the first inventory to be a multiple of your employee count. Deduplicate by credential, not by name.

Step 2 — Assign ownership. Every NHI gets a named human or team owner, an accountable cost center, and a purpose statement. Unowned identities get a defined quarantine path: restrict, observe,

retire.

Step 3 — Classify and right-size. Evaluate each identity's actual usage against its granted scope. Classical peer-group and usage analysis surfaces the over-privileged tail — typically the majority of the population.

Step 4 — Put lifecycle under policy. Creation through an approved path with an owner and an expiry; rotation on schedule; deactivation when the owning system or contract ends. The same reasoning engine that evaluates human access evaluates these policies continuously.

Step 5 — Watch context, not calendars. Re-evaluate when context changes: owner leaves the company, credential appears in a new system, usage pattern shifts, scope expands. Continuous reasoning replaces the annual campaign.

§ 04 — SECRETS & KEYS

The credential is the identity

For machine identities, the credential and the identity are inseparable: whoever holds the key is the account. Governance must therefore track credential inheritance — which pipeline, container, or engineer can read the secret that animates each identity. A vault tells you where a secret is stored; a reasoning layer tells you who can effectively use it, whether that inheritance violates policy, and when it should expire.

The question “which service account inherits this credential, and who owns it?” should have an automatic answer. In most enterprises today it is a two-week investigation.

§ 05 — GETTING STARTED

A ninety-day opening move

Days 1–30: inventory and ownership for the top three platforms by risk (typically cloud IAM, Active Directory, and the primary ERP). Days 31–60: policy evaluation live for creation and expiry on those platforms; quarantine process running for unowned accounts. Days 61–90: usage-based right-sizing recommendations flowing to owners; context-triggered re-evaluation replacing the next scheduled campaign for the covered population.

KeyForgeAI delivers each of these steps as part of its Non-Human Identity governance surface — one reasoning engine across employees, contractors, service accounts, secrets, and AI agents. See the companion paper, *Preparing for Agentic Identity*, for the next population arriving at scale.