

Non-human identity governance.

Service accounts, integration users and AI agents outnumber your people, and most have no owner, no review and no expiry. One identity model, an owner for every NHI, credentials you can date, a chain you can check link by link, and access you can take back.

RAPIDVALUE IGA · OCTOBER 2026 · www.rapidvalue.eu

01 The problem

NHIs are historically an afterthought: created by application teams without central registration, holding standing credentials that never rotate, over-permissioned, and unowned. AI agents make it worse. People connect them to the identity provider with one click, platforms let anyone build one, and they act on someone's behalf with nobody named. When the script breaks, or when a breach lands, nobody can say who is accountable or what the identity could reach.

We won't quote industry averages at you. The numbers that matter are the ones a first scan turns up in *your* estate. What follows is what we typically find:

PATTERN	TYPICAL FIRST-SCAN FINDING
Service accounts with no attributable owner	often half to three-quarters of the estate
Static credentials with no rotation on record	commonly years old; nobody dares touch them
NHIs holding admin-level access	a meaningful minority, rarely justified in writing
AI tools connected by a user's or admin's consent	present in most tenants, known to nobody
Attributing an NHI to a person mid-incident	days of archaeology, not minutes

02 One identity model for every NHI


RapidValue does not bolt a second system onto identity governance. A service account, an integration user and an AI agent are all identities, governed by the same ownership, review, reconciliation and audit machinery as people. Two declared facts decide how each is treated:

- **The lifecycle driver:** the asserted reason the NHI exists. It comes from an explicit assertion, an inherited account-type default, a visible conflict, a non-binding suggestion, or unknown. Only an asserted or

inherited driver may silence a driver-specific finding.

- **For AI agents, the agent profile:** a purpose, an autonomy tier (assisted, supervised or autonomous), the platform it lives on, and the person it acts for. A person's AI assistant that signs in to RapidValue is an agent identity of its own, with the person as its ceiling.

Agents are found where they live, through the connectors you already run: Microsoft Entra Agent ID, Amazon Bedrock and AgentCore, Google Vertex AI and Agentspace, Salesforce Agentforce, ServiceNow AI agents and declared GitHub Apps. Self-hosted agents are registered through an intake that says it is a declaration. Nothing becomes an agent because of its name.



One governed identity for every AI agent, as built in October 2026.

03 An owner for every NHI

Every NHI's owner is read through one resolver: an explicit owner, then a bucket, the system, the context, and finally a visible "no owner". A platform sponsor is imported as the owner fill-only: never over an owner a person set, and an unreadable sponsor gives no owner at all, so the NHI lands in the orphan inbox instead of getting a placeholder. Each owner records who asserted it and how.

The orphan inbox is where accountability is claimed: assign an owner, flag the NHI for decommission, or hand it to a stewardship group. Owners re-attest on a cadence: yearly for ordinary NHIs, and for agents every 90, 180 or 365 days by autonomy tier. The same cadence puts an agent's grants in an access review with its owner, who sees its purpose, tier, recent decisions and accepted terms before keeping or revoking.

04 Credentials you can date, never secrets you hold

A daily sweep inventories what each NHI holds: a reference, its kind and its dates. There is no column that could hold a secret. Findings come before the outage: expired, expiring soon, rotation overdue (a working credential older than 90 days), and no expiry at all. Last use is shown only when the source reports it; everywhere else it reads "not reported", never "never used".

sources today: **Entra** app secrets and certificates · **AWS** access keys · **Google Cloud** service-account keys · **Google Workspace** app passwords · **Salesforce** OAuth tokens of integration users · **ServiceNow** refresh tokens, API keys and OAuth registry entries · **Active Directory** password age · credentials RapidValue issued

05 Blast radius: what can this NHI reach?

Per NHI, RapidValue computes a four-level reachability graph: direct accounts and entitlements (L1), the systems reached through those accounts (L2), co-members on the same entitlements (L3), and the NHIs this identity owns (L4).

risk classification is **graph traversal**, not ML: deterministic and explainable to an auditor:
· ≥ 3 sensitive nodes or ≥ 4 systems → **critical** · ≥ 1 sensitive or ≥ 3 systems or blast>15 → **high**
· blast>7 → **medium**

06 The chain, link by link

For each NHI or agent, the chain view traces the owner, the person it acts for, its credentials, the tools it may use, the actions RapidValue recorded and the applications it can reach. Every link is marked **verified** (RapidValue read it from the system that holds it, or recorded it itself), **declared** (a person asserted it) or **unknown**, with the age of its evidence. A link nobody can vouch for is shown as unknown, never as fine. The same chain goes into a signed evidence pack.

07 Agents nobody registered

RapidValue reads the OAuth consents in Entra and Google Workspace every day and keeps the clients that look like AI agents or MCP clients as candidates, with a confidence and the names of the people who consented. An administrator adopts a candidate as a governed agent with an owner, dismisses it, or requests its revocation, which counts only when a complete read no longer sees the consent. A new high-confidence candidate raises a finding for the security team.

08 Every agent call decided

RapidValue's MCP gateway sits in front of your MCP servers and decides every tool call against the agent's grants, the person's rights and the owner's runtime terms. It starts in observe mode per system. In enforce mode an allowed call gets a single-use authorization and the system's credential is injected from the vault, so the agent never holds it. When the answer is "not yet", the gateway asks the owner (an access request, a JIT request or a confirmation of exactly that call) instead of refusing flatly, with a deadline that falls back to the IAM team. The gateway runs on staging today; production follows its load proof.

09 Taking access back, and proving it

Five triggers end an agent's access: the kill switch, a revoke in the access review, the end of a mission, the owner withdrawing a terms version, and stopping a person's assistant. All of them make the next call fail, because grants and status are read live; reconciliation then removes the access at the target. The kill switch

also disables the agent on Entra, AWS and Google Cloud where it can, and says plainly where it cannot. Releasing it never restores a credential that may have leaked.

Every step lands in the hash-chained audit log. Generated for one agent, the evidence pack is that agent's EU AI Act pack: terms, grants and approvals, human oversight, executions, incidents, retention, and an article map that says what it does and does not show.

10 The 90-day rollout

☐ Weeks 1–2: discover.

Connect AD and Entra, switch on the agent objects of the platforms you use, and run the first shadow-agent scan.

☐ Weeks 3–4: assert the reason.

Assert a lifecycle driver per account type, and a purpose and autonomy tier per agent. Unknown stays visible.

☐ Weeks 5–6: bulk-assign stewardship.

Per application team, so nothing stays in the global orphan pool. Triage the shadow agents.

☐ Weeks 7–8: claim campaign.

Each team claims its NHIs or flags them for decommission; ownership coverage is tracked as a KPI.

☐ Weeks 9–12: govern.

Attestations and agent access reviews live, credential findings worked, the kill switch proven on each platform, an evidence pack per critical agent.

THE KPI THAT MATTERS

Ownership coverage, the share of NHIs with an accountable, attested owner, is a top-line number from day one. DORA and NIS2 expect machine credentials to be governed like human accounts; this is how you show it.



Run the Quick Scan on your own directory. A POC surfaces your real NHI estate (count, owners, credentials, shadow agents, blast radius) in the first afternoon. www.rapidvalue.eu · hello@rapidvalue.eu