

WHITEPAPER

Preparing for *Agentic* Identity

AI agents request access autonomously, at machine speed. How to inventory agent lineage, validate declared scope, and enforce SoD on bots — before auditors ask.

Agentic access is not a future problem. If your enterprise is piloting AI agents anywhere — coding assistants, support bots, workflow agents, MCP-connected tools — those agents are already touching systems your governance program is accountable for. This paper is the pre-flight checklist.

§ 01 — WHAT CHANGES

Agents are not service accounts

Service accounts are static: fixed scope, fixed behavior, fixed integration. Agents are dynamic: they compose actions at run time, request access programmatically, and change behavior when their prompts, models, or tools change. An agent's effective capability is not what its credential grants — it is what its credential grants multiplied by what it can be instructed to do.

This breaks assumptions across the governance stack. Approval flows assume a human requester who can answer questions. Certifications assume the identity behaves the same next quarter. SoD engines assume access combinations accumulate slowly. None of these hold for agents.

§ 02 — FIVE QUESTIONS

The pre-flight checklist

Before agent traffic grows, your program should have automatic answers to five questions:

#	Question	Why it matters
1	Lineage — who built it, who deployed it, which human is accountable?	An agent without a traceable owner is a service-account problem wearing a newer costume.
2	Scope — does actual access match declared intent?	Scope drift happens faster in agents than any human population; requests arrive programmatically.
3	SoD — can it accumulate toxic access combinations?	Most SoD engines have never been pointed at bots. Combinations that would be flagged for a human sail through.
4	Expiry — does access die when the agent does?	Agents are forked, replaced, and decommissioned constantly. Their access lingers the way orphaned accounts always have.
5	Continuity — is it re-evaluated as context changes?	Agent context changes hourly. An annual review of agent access reviews a snapshot of something that no longer exists.

§ 03 — MCP & TOOLING

Governing the connective tissue

Model Context Protocol servers and similar tool gateways are becoming the interface between agents and enterprise systems. Each MCP server is, in governance terms, an entitlement bundle: it exposes scopes, and any agent connected to it inherits reach into whatever those scopes touch. Inventory them like applications: which scopes does this server expose, which are sensitive, which agents connect, and under whose authority?

Treat new server registrations and scope expansions as governance events that trigger reasoning — not as infrastructure changes that bypass it.

An agent's request should be evaluated with more context than a human's, not less — because the agent can neither explain itself in a hallway conversation nor be held accountable in a disciplinary one.

§ 04 — THE OPERATING MODEL

SoD and approval for things that act

Identity model: register every agent as a first-class identity with lineage metadata: owner, builder, deployment surface, model/tooling version, declared purpose and scope.

Request path: agent access requests flow through the same reasoning engine as human requests, evaluated against declared scope, SoD posture, data sensitivity, and the requesting agent's behavioral history.

Continuous validation: compare declared scope to observed behavior. Divergence triggers automatic narrowing or suspension, with the owner notified — the agent equivalent of a leaver process, executed at machine speed.

Audit posture: every agent decision explainable to the same standard as human decisions: named policy, named signals, named outcome. Regulators will not accept a lower bar because the requester was software.

§ 05 — GETTING AHEAD

Start while the population is small

The enterprises that struggled most with service-account sprawl were the ones that started governing after the population exploded. Agentic identity offers a rare second chance: most estates still count agents in the dozens. Register, own, and reason about them now, and the controls scale with the population instead of chasing it.

KeyForgeAI's Agentic Identity surface delivers this operating model — agent inventory, lineage tracking, scope validation, and SoD on bots — as part of one reasoning engine across human, machine, and AI identity. Assess your readiness at www.keyforge.ai/scorecard.