

WHITEPAPER

The Reasoning *Maturity* Model

A five-stage framework for scoring your identity governance program — from Manual to Agentic-Ready — with a practical roadmap for each stage.

Identity programs evolve along a predictable path. Knowing where yours sits — honestly — is the difference between an investment plan and a wish list. This paper defines the five stages of Reasoning Maturity, the observable behaviors of each, and what it takes to advance.

§ 01 — THE MODEL

Five stages of Reasoning Maturity

The model measures one thing: how much of your governance judgment is exercised automatically, continuously, and explainably — rather than by humans, episodically, and invisibly.

Stage	Name	Observable behavior
1	Manual	Spreadsheets, tickets, tribal knowledge. Access decisions depend on who you ask.
2	Platformed	IGA deployed, workflows running. Judgment still manual: approvers decide with little context; certifications are rubber-stamped.
3	Reasoned	Metadata-driven decisions. Risk-based routing live. Every signal evaluated automatically at request time; every outcome traceable to a policy.
4	Continuous	Non-human identities, service accounts, and secrets under the same reasoning as humans. Evaluation happens at every change in context, not at events.
5	Agentic-Ready	AI agents reasoned about continuously. Lineage tracked. SoD enforced on bots. Declared scopes validated against actual behavior.

§ 02 — DIAGNOSIS

Where programs actually sit

Most enterprises that purchased IGA between 2018 and 2023 sit at Stage 2. The tell-tale symptoms: certification campaigns with approval rates above 98 percent (reviewers are not reviewing); approver queues measured in hundreds; SoD checks that run in one ERP but nowhere else; and a service-account population nobody can enumerate.

Stage 2 is a ceiling, not a plateau. The platform cannot advance the program on its own because the missing capability — judgment — is not a workflow feature. Adding more workflows adds more queues; it does not add reasoning.

A program does not advance by buying more platform. It advances by automating the judgment the platform was never designed to exercise.

§ 03 — SCORING

How to score your program

Score each of the five governance surfaces (Access Governance, Lifecycle, Non-Human Identity, Risk & SoD, Agentic Identity) from 1 to 5 using the stage definitions above. Your program's stage is not the average — it is the *lowest* surface score that touches material risk. A program with Stage 4 access governance and Stage 1 non-human identity is a Stage 1 program from the auditor's chair.

Three questions calibrate each surface honestly: What fraction of decisions on this surface involve a human? Can you produce, for any decision, the named policy that produced it? How long after a change in context (role change, contract expiry, credential rotation) is the decision re-evaluated?

§ 04 — THE ROADMAP

Advancing stage by stage

From Stage 1 to 2 — platform the mechanics. Deploy or rationalize the IGA foundation: authoritative sources, connectors, joiner-mover-leaver automation, a request catalog. This is well-trodden ground; the industry knows how to do it.

From Stage 2 to 3 — automate the judgment. Introduce a reasoning layer above the platform. Codify the questions your approvers ask (or should ask) into policies evaluated automatically at request, approval, and review: privilege classification, training status, SoD posture, JIT eligibility, expiry alignment. Target the top ten decision types by volume first; they typically cover more than 80 percent of tickets.

From Stage 3 to 4 — extend to everything non-human. Inventory service accounts, API keys, secrets, and machine identities. Assign ownership. Put them under the same policy evaluation as humans, continuously — not on a quarterly campaign.

From Stage 4 to 5 — govern the agents. Track agent lineage (who built it, who deployed it, who is accountable). Validate declared scope against actual behavior. Enforce SoD on bot access combinations. Expire agent access when agents are decommissioned or forked.

§ 05 — MEASUREMENT

KPIs that prove progress

Maturity claims need instrumentation. The KPIs that matter:

- **Automation rate:** percentage of access decisions completed with zero human touches, by surface.
- **Explainability rate:** percentage of decisions traceable to a named policy on demand.
- **Context latency:** time from a change in context to re-evaluation of affected access.
- **NHI coverage:** percentage of non-human identities with a named owner and lifecycle policy.
- **Certification signal:** revocation rate in reviews (a 0.5% revocation rate means the review is theater).

§ 06 — NEXT STEP

Score yourself in five minutes

The interactive version of this model — fifteen questions across the five surfaces, scored automatically with a gap report mapped to your IGA platform — is available at www.keyforge.ai/scorecard.