

DATASHEET

The Three-Tier *Reasoning* Engine

Deterministic, statistical, and generative reasoning — each independently deployable, each fully transparent. You choose where to stop.

OVERVIEW

Reasoning, defined

We define reasoning the way computer science defines it: the systematic evaluation of multiple signals to reach a defensible conclusion. KeyForgeAI's engine operates in three tiers. Regulated buyers can stop at Tier 1 and run a complete, fully deterministic engine with no generative AI in the decision path. Innovation-led organizations can run all three. Every decision in every tier is explainable, auditable, and traceable to a named policy.

TIER 01 — DETERMINISTIC

The Policy Foundation

Rule-based evaluation of metadata at every decision point. *Example: if the entitlement is privileged AND the requester is a contractor AND the contract expires within 30 days, route to a compliance approver and set expiry to contract end date.* Every outcome is traceable to a named rule. Zero opacity.

Attribute	Value
Technology	Rules · Policies · Metadata
Audit posture	Fully deterministic
Availability	Required for all deployments

TIER 02 — STATISTICAL

The Pattern Layer

Peer access analysis, role mining, outlier detection, usage-based right-sizing. Classical machine learning — clustering, classification, regression — industry-standard in IGA for over a decade. Surfaces patterns humans cannot see manually; fully explainable through feature importance.

Attribute	Value
Technology	Classical ML · Statistics
Audit posture	Explainable models
Availability	Optional; recommended for most

TIER 03 — GENERATIVE

The Future Layer

LLM-assisted policy authoring, natural-language access requests, AI-agent intent analysis, conversational audit explanation. Opt-in and off by default — never inserted into the decision path without explicit configuration.

Attribute	Value
Technology	LLMs · Generative AI
Audit posture	Opt-in, isolated
Availability	For agentic-ready organizations

Run Tier 1 alone for a fully deterministic engine with no generative AI in the decision path. Add Tier 2 for pattern intelligence. Enable Tier 3 when your organization is ready. The architecture meets you where you are, and grows with you.

COMMON QUESTIONS

What auditors and architects ask

Can we prove no LLM touched a decision? Yes. Tier 3 is architecturally isolated; when disabled, no generative component exists in the decision path, and the evidence trail shows it.

Can tiers be enabled per surface? Yes. Many customers run Tier 2 for access governance while keeping non-human identity on Tier 1 only.

What happens when policies conflict? Deterministic precedence rules resolve conflicts, and the resolution itself is recorded in the decision evidence.

NEXT STEP

Evaluate against your controls

Request the architecture briefing at www.keyforge.ai/contact — a working session with your identity and security architects on tier configuration for your regulatory posture.