



DATA SHEET

Agent IQ

Measure every agent run, prove every workflow's ROI, and govern autonomous execution — before the cost hits your invoice.

HEAD OF AI

CFO

VP ENGINEERING

THE CHALLENGE

Agentic AI costs accumulate silently. Agents run in the background across teams and workflows — charges surface on the cloud invoice weeks later with no workflow-level breakdown.

There's no standard way to measure agent ROI. Every team defines success differently. Without a shared measurement layer, the CFO can't evaluate agent spend and the Head of AI can't compare deployments.

Governance gaps widen as agents scale. Autonomous execution means errors propagate — and policy violations complete — before anyone in the loop can intervene.

What Agent IQ Does

Agent IQ gives you real-time observability into every agentic AI workflow your organization runs — autonomous agents, multi-step automations, and custom-built orchestrations. Every execution is captured at the platform layer: cost, latency, success/failure, step-level trace, and the business KPI it was supposed to move.

On the governance side, Agent IQ evaluates every agent run against your policies before the workflow completes. Anomalous behavior, unexpected data access, and out-of-bounds execution attempts are flagged and logged — giving your CISO and Head of AI a control surface across deployments they didn't build and can't directly observe.

KEY CAPABILITIES

● Execution analytics

Success rate, latency, step-level trace, and failure mode distribution for every agent run — by workflow, agent, and team.

● Cost-per-workflow measurement

What each agent execution actually costs — across LLM tokens, tool calls, and compute — attributed by workflow, vendor, and cost center.

● Custom KPI tracking

Connect agent runs to the business outcomes they're supposed to produce. Define KPIs per workflow; Agent IQ scores every execution against them.

● Agent comparison

Head-to-head ROI, cost, and success rates across agents solving the same problem. Know which to standardize, which to retire.

● Governance and policy rails

Flag anomalous runs, surface failure patterns, and enforce execution guardrails — across both managed SaaS agents and custom-built workflows.

● Kai synthesis

Ask in plain English: "Which agent workflows are ROI-positive and which are burning budget?" Kai synthesizes Agent IQ data and answers with reasoning shown.

HOW IT WORKS

1

Connect your agent infrastructure

API, SDK, or webhook integration in minutes. Agent IQ begins observing every workflow execution immediately — no instrumentation of individual agents required.

2

Every run is auto-analyzed

- Cost, latency, success/failure, KPI contribution, and policy compliance scored on every execution. Step-level trace available for failure investigation.

3

Measure, compare, decide

- ROI dashboard by workflow, agent, and team. Flag anomalous runs. Scale what's working. Retire what isn't. Brief the board in seconds via Kai.

METRICS YOU'LL TRACK

Metric	What It Tells You
Execution success rate	% of runs completing without failure, by workflow
Cost per execution	Spend per agent run — by workflow, vendor, and team
Latency by step	Time at each stage of a multi-step workflow
ROI per workflow	Business value generated relative to execution cost
Failure mode distribution	Error types by frequency and business impact

Metric	What It Tells You
Workflow KPI score	Custom business outcomes achieved per run
Agent comparison score	Head-to-head ROI across agents on the same task
Anomaly rate	Runs that deviated from expected behavior
Team utilization	Which teams run agents and at what volume
Governance flags	Policy violations and out-of-bounds execution attempts

WORKS WITH

Devin · Cursor (agentic mode) · GitHub Copilot Workspace · Claude Code · AutoGPT · LangChain · n8n · Zapier AI · AWS Bedrock Agents · Google Vertex AI Agents · Microsoft Copilot Studio · and more

5×

agentic AI spend growth YoY — the fastest-growing cost center with zero visibility

0

lines of agent instrumentation required — observes at the platform layer

100%

of agent runs captured — approved, custom-built, and unsanctioned alike

See Agent IQ in action.

Drop your work email at olakai.ai/agent-iq/ — Magic Link, 7-day access, no demo call required.

olakai.ai/agent-iq/

Magic Link · 7-day access · no account needed