Relevance AI

DATASHEET

RELEVANCE AI FOR CIOS

A future-proof, agile platform for enterprise AI agents

Business experts build and run governed AI agents on one secure, scalable platform. IT can set guardrails and keep central oversight, without building or maintaining the agent stack, or worrying about model updates and swaps.

> **“I don’t want to be the single author of any agent.”**
>
> Rakuten
>
> **Benjamin Cox**
> VP of AI Strategy & Operations, Rakuten Advertising
>
> **Rakuten Advertising** runs a hub-and-spoke model: a core team handles major builds, analysts embedded in business units build where it matters, and governance stays on the big projects.

Building agents is easy. Running them in production is hard.

Agents stood up one at a time tend to become one-off projects, each with its own integrations, data, and tooling. Whenever requirements change, they need to be rebuilt. Agent pilots often stall on what it takes to truly make an agent workflow enterprise class and production ready: governance, monitoring, evaluations, and security sign-off.

Own your business logic, not the agent infrastructure

Relevance ships the agent stack as one governed system: context layer, MCP gateway, LLM router, evaluations, orchestration, job queue, triggers, monitoring, and tracing. Your engineers stay focused on the work that sets your business apart, and the agent you pilot is the agent you run in production.

Central governance, business-led building

Domain experts build agents and manage their performance. IT keeps a single pane for access control, permissions, and data security, with approval gates and version control.

Enterprise security and compliance

SOC 2 Type II and GDPR. Encryption in transit and at rest, SSO/SAML, RBAC, audit logs, PII masking, data residency, and no training on your data.

Live in weeks, with a light lift for IT

An embedded deployment team maps use cases in weeks 1–2 and deploys your first team of pilot agents from week 3, then trains your team to build on its own.

TRUSTED BY TEAMS AT

Canva · databricks · CONFLUENT · AUTODESK · lightspeed · Rakuten · freshworks · AVEVA · employmenthero · QUALIFIED

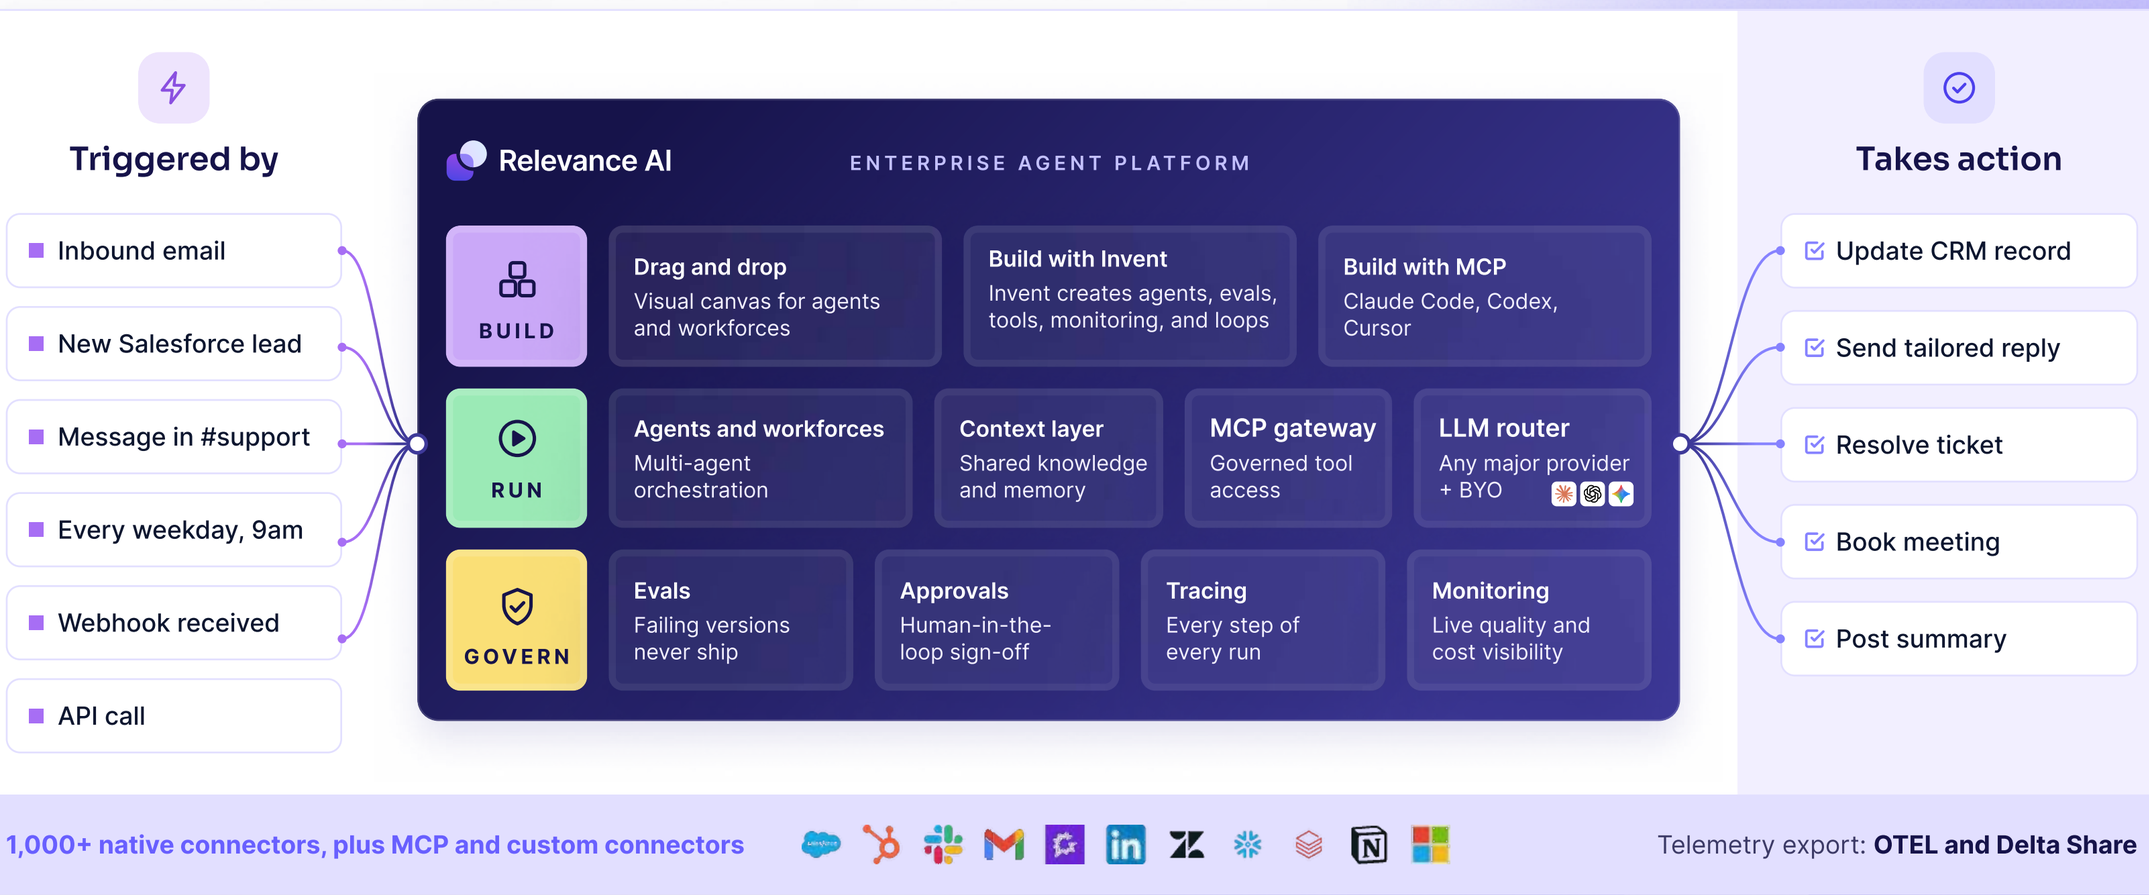
One governed platform for business-critical agents

Trade chaos for calm. All key agents on one stack, instead of a patchwork for IT to maintain

Most teams bolt together a router, a queue, an eval tool and a tracer just to run agents. Relevance ships it all as one system, so teams can build quickly and easily, you can catch regressions before they ship, run each agent on the most affordable model that meets your quality bar, score live runs (not just test sets), govern tool access in one place, retry failed runs instead of losing them, and see what every task really cost.

Business teams build, IT sets the guardrails

Three ways to build: drag and drop on a visual canvas, describe the agent in plain language and let Invent build it, or drive the platform over MCP from Claude Code, Codex, or Cursor. Every path ships to the same runtime, gated by the same evals.



1,000+ native connectors, plus MCP and custom connectors

Telemetry export: **OTEL and Delta Share**

Built to satisfy IT requirements

SECURITY

Enterprise-grade security

- **Data protection:** AES-256 encryption in transit, encryption at rest, PII masking, and no training on your data.
- **Data residency:** hosted on AWS in the US, EU, and Australia.
- **Access and audit:** SSO/SAML, role-based access control, audit logs, and full agent tracing.

SOC 2 Type II and GDPR. Audit reports and policies at trust.relevanceai.com

CENTRALIZED GOVERNANCE

One pane of control

- **Access:** a single pane for access control, permissions, and data security, with per-tool approval rules and governed tool access.
- **Lifecycle:** draft, active, and pinned versions. Every change re-runs its eval suite, and a failing version never ships.
- **Oversight:** high-stakes actions pause for human sign-off at thresholds you set. Admins can switch Invent on or off per project.

SCALABILITY

Runs at enterprise scale

- **Managed queue:** when thousands of tasks land at once, the queue paces the work so that nothing downstream falls over.
- **Built-in resilience:** autoscaling workers, concurrency caps, retries with backoff, a dead-letter queue, and work-hours pacing.
- **Parallel work:** parallel and branching runs, plus bulk runs across datasets.

EASE OF INTEGRATION

Fits your existing stack

- **Connectors:** 1,000+ native connectors, plus MCP support and the ability to build custom connectors.
- **Triggers:** native triggers for CRMs, email, calendar and more, plus webhooks, cron, custom triggers, and API calls.
- **Observability:** export telemetry through OTEL and Delta Share.

MINIMAL IMPACT ON IT

Less for IT to build and maintain

- **No stack to assemble:** one system instead of a separate router, queue, eval tool and tracer.
- **Experts own quality:** domain experts can set quality standards, define guardrails and monitor every agent, not only engineers.
- **Invent does the build work:** The Invent solution engineering agent in the Relevance platform can create agents, tools, evals and monitoring from a plain-language brief.

FUTURE-PROOFING

Model-agnostic by design

- **A portfolio of models:** route to any major provider, including Claude, GPT, Gemini, GLM and Kimi models, or bring your own. No lock-in.
- **Best ROI for each job:** evals measure performance; the lowest-cost model that clears your performance standard wins.
- **Low-impact model updates:** when a new model ships or updates, the agent's eval suite is re-run *before* publishing. Fallback model on failure.

Build or buy? A complete platform simplifies complexity

A coding agent can write the basic agent loop in an afternoon. The platform around it is a bigger challenge and key to enterprise-class scalability, reliability, and security: the identity layer, integration layer, runtime, memory system, governance model, evaluation loop, and interface. Every piece you decide to build on your own becomes infrastructure that your team has to upgrade, monitor and maintain. Success with agent platforms is only half about technology. The other half is knowing *how* to deploy agents inside a real enterprise.

The easy part: call a model, let it choose a tool, run the tool, feed the result back, repeat.

The hard part: the eight layers below, built, secured and maintained for every agent in the company.

Models and routing

IF YOU BUILD IT

Multiple model vendors, a router or model proxy, fallbacks, and a designated owner for the proxy.

WITH RELEVANCE

LLM router for any major provider plus bring-your-own, with a fallback model on failure.

Integrations and identity

IF YOU BUILD IT

Per-user OAuth, token refresh and revocation, plus the custom tools no MCP marketplace covers.

WITH RELEVANCE

1,000+ native connectors, an MCP gateway, custom connectors, and sandboxed Python and JavaScript tools.

Triggers

IF YOU BUILD IT

Webhooks, polling, payload mapping, deduplication, retries, replay, and loop prevention.

WITH RELEVANCE

Native triggers, webhooks, recurring schedules, 100s of pre-built events, custom signals, and API calls.

Runtime and queue

IF YOU BUILD IT

Durable jobs, fair scheduling across teams, a dead-letter queue, and per-vendor rate limiting.

WITH RELEVANCE

Managed queue with autoscaling workers, concurrency caps, retries with backoff, and work-hours pacing.

Context and memory

IF YOU BUILD IT

Continuous sync from Drive, Notion, SharePoint and Confluence, semantic search, summarisation, and memory.

WITH RELEVANCE

One shared context layer with hybrid RAG, automatic syncing from key data sources, and long-term memory.

Orchestration and approvals

IF YOU BUILD IT

Shared scratchpads, conditions, escalation paths, A2A orchestration, memory, and decision models.

WITH RELEVANCE

Multi-agent graphs with handoffs, parallel and branching runs, and conditional approval rules.

Governance and quality

IF YOU BUILD IT

Drafts, versioning, roles and SSO, plus eval suites, publish gates, and production scoring.

WITH RELEVANCE

RBAC, SSO/SAML, draft, active and pinned versions, pre-publish evals, and production sampling.

Interfaces and dashboards

IF YOU BUILD IT

Chat, Slack, Teams, mobile access, your own MCP server, plus logs, cost attribution, and activity reporting.

WITH RELEVANCE

Embeddable chat, voice agents, API and MCP, full agent tracing, cost visibility, and OTEL export.

Easy to get agents in production, with proven results.

INVENT

The agent that builds the agents, so business teams don't have to wait on IT

Domain experts collaborate with Invent to define the process and the desired outcomes. Invent drafts the agents, evals, and monitoring, and can make recommendations for fixes in the future. Expertise carries forward through model upgrades and vendor changes.

1 Create

Learns your process from a brief, connects your systems, and staffs a team of specialist agents.

2 Evaluate

Sets the checks every agent must pass. Agents go live only after they pass the tests you require.

3 Monitor

Samples live runs against the bar you set and flags the moment quality drifts.

4 Improve

Proposes the fix, runs the checks against it, and suggests new agents worth building.

TECHNOLOGY

QUALIFIED

Qualified generated $7M pipeline with 35+ AI agents across the org

$7M Pipeline in 6 months

10x Increase in output

35+ Agents across the org

Chose Relevance AI after evaluating 60 vendors. The first BDR agent performed at human-level quality in one week.

HOW IT RUNS IT

Canva

Business teams build on governed platforms. IT sets the guardrails.

Canva's GTM and IT teams start in sandbox environments with human-in-the-loop review. Workflows move into production only once governance, monitoring, and user experience are in place, and then ownership shifts to the business team.

Recognized by the industry

CBINSIGHTS
Leading Enterprise Agent Vendor

Everest Group®
Agentic AI Luminary

Capgemini
Multi-Agent System Leader

4.5 stars on G2