




SERAA Cortex · From first build to governed production

Every agent. Every cloud. One control plane.

Build, orchestrate, register and monitor AI agents, including the ones you didn't build, under one governance layer across the clouds you already use.

ISO 42001 · SOC 2 · ISO 27001

Sovereign and air-gapped

GKE · AKS · EKS · GCP Agent Engine

Covasant is one of nine partners selected globally by Google Cloud to build agentic AI practices on Gemini Enterprise.

Agents multiply faster than anyone can count them.

TODAY

"How many AI agents are we running?"

Four teams, four clouds, and nobody can produce a number.



WITH AGENT REGISTRY

One catalogue, auto-discovered from Gemini Enterprise, Azure AI Foundry, AWS Bedrock and Databricks AgentBricks.

TODAY

"Which version is live, and who approved it?"

The answer is in someone's inbox.



WITH AGENT REGISTRY

Gated promotion from Dev to Production, with semantic versioning, version diffing and full lineage.

TODAY

"What did that agent cost last month?"

Four invoices, none tied to a use case.



WITH AGENT CONTROL TOWER

Cost ranked by agent, model, user, provider and tenant, logged on every call.

TODAY

"Did it stay inside the rules?"

You find out after the incident.



WITH AGENT CONTROL TOWER

Runtime guardrails for content, PII, GDPR and HIPAA, applied whatever the model.

The structural point: every platform governs its own

Hyperscalers, model providers and agent builders each see what runs on their own platform, and see it well. No cloud can sell governance across the others without arguing against its own consolidation pitch. SERAA Cortex sits above all of them.

From first build to governed production.

Four connected modules (Build, Orchestrate, Register and Monitor) carry every AI agent from first prompt to live production, each handing off to the next. Wrapping every stage, Assure, the governance layer, keeps each agent controlled, policy-bound and auditable from day one.



Any model. 100+ LLMs across Google, Anthropic, OpenAI, Azure, AWS and Databricks.

Open protocols. A2A for agents working together, MCP for tools.

Any runtime. GKE, AKS, EKS or GCP Agent Engine, including air-gapped.

Four modules, each useful alone.

BUILD

Agent Work Bench

Build agents your way.

- ✓ No-code and pro-code agents in one interface; pro-code scaffolds on LangGraph
- ✓ Model Hub: 100+ LLMs across six providers
- ✓ MCP Hub: tools registered centrally with scope, inputs and outputs
- ✓ Knowledge bases over Vertex AI Search, Azure AI Search, OpenSearch, FAISS and more

ORCHESTRATE

Multi-Agent Orchestration

Compose agents into work.

- ✓ Visual canvas for sequential and parallel branch-and-merge flows
- ✓ A2A protocol support: third-party agents compose natively
- ✓ A different model at each stage, so cost is set at design time
- ✓ Draft, then publish to QA, Staging or Production

REGISTER

Agent Registry

One catalogue for every agent.

- ✓ Auto-discovery from Gemini Enterprise, Azure AI Foundry, AWS Bedrock and Databricks AgentBricks
- ✓ Gated promotion from Dev to Production, with full lineage
- ✓ Semantic versioning, with version diffing before publish
- ✓ Deploy to GKE, AKS, EKS or GCP Agent Engine from one workflow

MONITOR

Agent Control Tower

See and govern every call.

- ✓ Searchable log of every LLM call: session, model, tokens, per-call cost
- ✓ Cost ranked by agent, model, user, provider and tenant, with budget caps
- ✓ Runtime guardrails: content filters, PII redaction, GDPR and HIPAA
- ✓ Monitors agents on other platforms when their calls route through SERAA Cortex

Agent Test Bench: evidence for the sign-off meeting.

Tests an AI agent or agentic workflow against a simulation built from your own historical logs, with expert, novice and adversarial personas. Execution validation against systems like SAP or Oracle is on the roadmap.



One control plane over every agent.

Agent Registry and Agent Control Tower each stand alone, so you can bring governance and live oversight to agents built on Gemini Enterprise, Azure AI Foundry, AWS Bedrock or Databricks AgentBricks, without moving or rebuilding one.



On agents built elsewhere. Agent Control Tower monitors their model calls, cost and guardrails when those calls route through SERAA Cortex as the model gateway and telemetry is configured. An agent that calls models directly is catalogued, but not observed at call level.

SCOPE-BASED ACCESS CONTROL

Assets flow up. Infrastructure flows down.



Business units build on their own while governance sees everything beneath it. A project can never use a model its parent has not approved. One account can be a Viewer at System level, an Admin in its Organization and a Contributor on a single Project, with SSO through Microsoft Entra ID, GCP IAM and AWS Cognito.

Assessing your agents.

Four things about any agent in production: **what it did, what it cost, whether it was right, and who answers for it.** Most enterprises can answer the first two on one platform. SERAA Cortex answers all four.

INVENTORY 1 Which agents are running, where, and who owns them? Agent Registry One catalogue, auto-discovered from four vendor platforms, each agent with an owner.	ACTIVITY 2 What did this agent do in this session? Agent Control Tower A searchable log of every model call, plus live status and success rate per agent.	COST 3 What did it cost, and who pays? Agent Control Tower Cost by agent, model, user, provider and tenant, with budget caps and alerts.
QUALITY 4 Was the answer right, and is it still right? Agent Test Bench Scored for faithfulness, context precision and task success, re-scored to detect drift.	POLICY 5 Did it stay inside the rules? Agent Control Tower Runtime guardrails: content filters, PII redaction, GDPR and HIPAA policies.	ACCOUNTABILITY 6 Who changed it, who approved it, and can we prove it? Agent Registry Gated promotion, versioning and lineage, and a log of every platform action.

Application monitoring was not built for agents

	Application monitoring	Agent observability
Unit of work	A request to a service	A session: many model calls, tool calls and decisions
What failure looks like	An error or a timeout	A confident wrong answer, an unapproved action, a cost spike
Accountability	The on-call engineer	A named owner, a version, an approver

Observe at the gateway. Every model and tool call routed through SERAA Cortex passes one controlled point; an agent that calls models directly is catalogued but not observed at call level. The result is a complete, searchable record of every routed call, with PII detected and redacted at runtime.

Tokens are only the first layer.

Consumption pricing, multi-step agents and multi-vendor estates have made AI a variable cost. One request can trigger many model calls, and the bill arrives after the work is done.



Roadmap

Choosing what happens at a limit (alert, switch to a secondary model, or stop), recommendations from your own traces, and a suggested mix of models are in development. Direction, not a commitment.

Where to start

Start with one agent, or all of them.

70%

faster time-to-market for new agents*

40–60%

lower LLM cost through governed routing*

100+

LLMs available in the Model Hub

Four ways in

Agent sprawl, and no inventory

Agent Registry, over the agents you already run

No view of AI spend or policy

Agent Control Tower, with calls routed through its LLM Gateway

A first agent to build, governed from day one

Agent Work Bench and Multi-Agent Orchestration

An agent nobody will sign off for production

Agent Test Bench, on your own historical logs

BUILT FOR REGULATED ENTERPRISES

ISO 42001 certified

SOC 2

ISO 27001

Aligned to GDPR, HIPAA, EU AI Act

SSO: Entra ID · GCP IAM · AWS Cognito


Per-customer VPC

Sovereign and air-gapped

Full LLM call audit log

Map your agent estate against the six questions on page 6.

Scan to book a briefing with our architects and see where the agents you run stand today.



[Request a Briefing](#)

Every agent. Every cloud. One control plane.

SERAA Cortex is the Act layer of SERAA, alongside SERAA Pulze for revenue signal, SERAA Axon for data intelligence and reasoning, and SERAA Synapse (alpha) for governed enterprise context.

London, UK · Plano, TX (USA) · Hyderabad, India · Dubai, UAE

* Figures reported by Covasant from product benchmarks and customer engagements. Not audited averages. We will show the workings on request.