

SOLUTION BRIEF

Agentic AI Accelerator

One Planner, Many AI Agents

One Planner agent, and specialized AI agents for SQL, APIs, documents, telemetry, and reporting — turning operational data into answers and automated action.

YOU OWN THE IP · NO PRODUCT COSTS · NO ROYALTIES



VividCloud's Agentic AI Accelerator is our proprietary framework for building you a modular, AWS-native agentic system — one central **Planner agent** and a team of specialized AI agents for operational data, external APIs, documents, and reporting.

The Planner reads each request and decides, step by step, which specialist agent to call next — querying data, calling external systems, searching documents, and building visualizations — before composing a single, cited answer, streamed back visibly the whole way.

Domain configurations swap in without touching agent code, letting one shared engine specialize by vertical — from robotics and warehouse automation to property & casualty insurance.

Key Features



One planner, looping until it's sure

Planner-Orchestrated Intelligence:

A central Planner agent evaluates every request, decides which specialist to call next, and feeds each result into shared context — looping up to five steps until the answer is complete.



A team of specialists, not one model

Purpose-Built Agents:

Dedicated agents for operational data (writes, executes, and self-repairs queries against events, telemetry, and logs), external APIs (auto-discovers OpenAPI specs, plans and repairs calls), document knowledge (hybrid vector + keyword search, reranked, cited, and automatically re-searched when the first pass falls short), and visualization (drafts, renders, and redraws charts until they actually communicate the answer).



Ask it, or wire it in

Two Ways In — Conversational & Headless:

Conversational — for people who need to dig in: analysts and adjusters get an interactive chat experience that streams reasoning as it happens, not just a final answer. **Headless** — for systems that need an answer: authenticated services call `POST /query/stream` and get the same event-streamed reasoning with zero UI involved.



One engine, tuned per vertical

Built to Specialize:

Domain configurations swap in without modifying agent code, letting the same Accelerator-built system be tuned for different industries while every specialist agent stays reusable underneath.



Nothing operates as a black box

Continuous Memory & Explainable Results:

The Continuous Memory agent extracts durable facts from every conversation and queues them for expert review before they inform the next answer, while every response carries provenance and reasoning context — auditable and reproducible for enterprise and compliance teams.



A secure, AWS-native foundation

Enterprise-Grade Infrastructure:

Built on Amazon Bedrock, Cognito, Aurora with pgvector, Elasticsearch/OpenSearch, and DynamoDB, with ECS/Fargate plus Lambda provisioned via CDK — for fine-grained access control and consistent, repeatable deployment.

Where It's Live Today

Robotics & Warehouse Automation

A floor supervisor asks about a downtime spike and gets a cited answer in seconds. A fleet-monitoring service asks the same question headlessly every 15 minutes, auto-triggering a reroute or maintenance ticket.

Property & Casualty Insurance

Adjusters cross-reference claims against loss history and cited guidelines, while a claims-intake system calls the same agents on First Notice of Loss to route standard processing versus fraud review.

One engine. New verticals swap in as a configuration change, not a rebuild.

Let's Talk

For technical architecture, configurability, security, and speed, see the FAQ and related documents for this offering.

150 Admiral Fitch Ave, Brunswick,
ME 04011
100 Domain Drive, Exeter, NH
03833
(781) 645-7800

VIVID
CLOUD