


Auraa

The Fastest Path to AI-Ready Data on Databricks

Auraa delivers governed, AI-ready Databricks lakehouses faster and at significantly lower cost than traditional implementations. Auraa, our proprietary agentic delivery platform, ensures every engagement is automated, governed, and production-ready from the first day on site.

Tier 1 consulting firm outcomes. Without the headcount. Without the extended timelines.



4 hours

To deploy Auraa

< 15 min

First Data source live, no code

70% lower

Total program cost

77% smaller

Team vs. traditional SI

The Problem

Databricks customers across enterprise wait many months after purchase, writing pipeline code, onboarding sources manually, chasing quality issues, and fighting maintenance that never ends. Covasant closes that gap using Auraa, our agentic delivery platform. The customer receives the outcome of a governed, AI-ready lakehouse.

01

The Talent Gap

2.7 M unfilled data engineering roles in the US. Mid-market and enterprise organizations alike struggle to hire, retain, and deploy data engineering talent at the pace their AI agenda demands.

Auraa delivers outcomes at Tier 1 consulting firm quality, without the headcount or the extended timeline. Auraa is the engineering capacity you need, at any scale.

02

The Backlog That Never Clears

Every organization has a queue of data sources the business wants onboarded. It never clears, blocking every downstream AI initiative.

Auraa clears the backlog. Any source waiting for months can be live in days, governed, tested, and production-ready.

03

The Activation Gap

Databricks is purchased. Databricks activation stays flat. The AI use case that justified the investment keeps moving further out.

Auraa drives higher Databricks value from day one, turning a stalled platform into a producing asset.



What We Deliver And How

Covasant uses Auraa, our proprietary agentic platform built exclusively on Databricks to automate the routine 70% of data engineering work. You receive a production-grade, governed lakehouse. Here is exactly what gets activated in your environment, and the workloads it unlocks.

Databricks Stack We Activate

- **Unity Catalog** - automated provisioning, tenant isolation, full lineage tracking and RBAC from day one.
- **Delta Lake** - all data stored natively in Delta; Bronze, Silver, and Gold layers managed end-to-end.
- **Lakeflow Declarative Pipelines** - driven by Auraa as the execution layer for automated ingestion.
- **Databricks Runtime** - all compute runs in your workspace; no separate environment, no data leaves your platform.
- **Genie Code via MCP** - 166+ tools exposed through 3 MCP slots; Auraa provides the context Genie performs the operations.
- **MLflow** - model tracking and experiment management active from first AI workload deployment.

Workloads We Drive

- **Data ingestion at scale** - multi-source Bronze layer automation across JDBC, SAP, flat files, and streaming sources.
- **Pipeline automation** - Bronze-to-Silver transformation built and governed without manual pipeline code.
- **Data quality enforcement** - profiling, rule generation, drift detection, and cross-table DQ active on every ingested dataset.
- **Governance workloads** - Unity Catalog registrations, access policy management, and full audit trail generation.
- **Gold layer transforms** - canonical Silver-to-Gold entity modelling; the foundation every AI use case consumes.
- **AI readiness** - trusted, lineage-tracked, governed data in the Gold layer - ready for models, dashboards, and agents from day one.

How An Engagement Works - Four Steps

No multi-month implementation program. No large consulting team. Auraa deploys, connects your sources, and hands you a governed lakehouse - in weeks.

Step 01

Deploy

Day 1 - 4 hours

Auraa provisions your Databricks workspace, configures Unity Catalog, and initialises governance controls. Production-grade before end of day.

Step 02

Connect Sources

Day 1 - 7

Data sources are onboarded conversationally - no engineering tickets, no schema design sprints. Bronze layer pipelines build and validate automatically.

Step 03

Data Flows - Governed

Week 1 onwards

Bronze-to-Silver processing runs autonomously. Quality rules enforced on every record. Lineage tracked end-to-end. Schema drift handled without engineer intervention.

Step 04

AI Use Cases Unblocked

Weeks, not months

Trusted, governed data in the Gold layer. Your team builds AI products, models, and analytical applications on a foundation that holds - not one that needs constant maintenance.

Platform Timeline - What's Available When

Auraa's capability evolves across three horizons. Every engagement today leverages the Foundation layer. Post-delivery, customers retain Healix, available in October 2026, to manage and govern what was delivered.

Now - June 2026

Foundation

Live

- Deploy in 4 hours
- Source onboarding in < 15 min
- Bronze-to-Silver automated
- Unity Catalog governance active
- Schema drift handling
- Genie Code MCP - 166+ tools

Pilot programme: 4 spots, 90-day commitment.

Aug 2026

Enterprise Scale

In Delivery

- Full Gold layer (Bronze → Silver → Gold)
- Runtime governance enforcement
- Cross-table data quality (DCC)
- Advanced RBAC + SSO + CLI
- Full audit trail + end-to-end lineage

Engagements open – Retail, Consumer Goods, Manufacturing, and Tech.

Oct - Dec 2026

Agentic Maturity

Roadmap

- Multi-agent workflow orchestration
- Autonomous end-to-end delivery
- AI decision explainability
- SAP + Salesforce connectors
- SOC 2 Type II certification

Zero manual steps from source discovery to Gold layer.

Oct 2026 onwards

Healix - Post-Delivery

Post-Delivery Layer

- Manage and govern what was delivered
- Operational monitoring and alerting
- Self-healing for delivered pipelines
- Governance continuity - no new sources

Customers retain Healix after engagement. New sources or pipelines require a new Covasant engagement.

Auraa Vs Traditional Build

Dimension	Traditional SI / Build	Auraa - Led Delivery
Time to first value	Many months to production	4 hours to deploy · first source live in under 15 minutes
Team size	Army of engineers required	77% team size reduction
Total cost	Premium SI programme cost	70% lower total programme cost
Source onboarding	Weeks per source - engineering tickets, manual design	Under 15 minutes - conversational, no code, no ticket
Governance	Manual, fragmented - added late or omitted	Automated from Day 1 - every decision auditable
Pipeline maintenance	Engineer on-call for every failure	Schema drift handled autonomously - no pages, no tickets
Scalability	Linear with headcount	Elastic with Databricks compute - no team growth required

Limited Pilot Programme - 4 Spots Available

Pricing Preferential Pilot program only	Commitment 90 days No long-term obligation	Cohort 4 customers Intentionally limited	Status Now open Spots allocated in order
---	--	--	--

30 Minutes. A Clear Picture Of What's Possible.

Schedule Your 30-Min Discussion

We run simulation for your exact business problems, show you which stalled sources can be live this quarter, and answer every question your team has - honestly. No sales scripts. No pressure.

About Covasant Technologies

Covasant Technologies is an AI-native Services-as-Software company, pioneering the category where Agentic AI elevates enterprise operations by shifting routine workflows to intelligent, autonomous systems, allowing human teams to focus on high-value strategy and oversight. At its core are two flagship platforms, **Covasant Agent Management Suite**, a comprehensive platform to build, deploy, and scale AI agents across any business function and industry; and **ARIA**, a core cognitive engine that unifies enterprise data into the advanced reasoning intelligence that powers these agents. Backed by deep expertise in data engineering, AI engineering, cloud infrastructure, and enterprise risk management, Covasant helps global enterprises transition from AI ambition to live production within weeks, not quarters, turning AI investments into measurable business outcomes with the reliability, scale, and control that modern enterprises demand.



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