

ARVIX AI: NOT AN ADVISOR. AN OPERATOR.

THE AUTONOMOUS OPTIMIZATION ENGINE

An LLM starts from zero. Arvix AI never does. It's already analyzed 10B+ workloads across 100+ enterprises and knows what fails at your scale.

10B+

WORKLOADS ANALYZED

100+

ENTERPRISES

5.95x

AVERAGE YEAR-ONE ROI

The Challenge

Generic AI tools and point solutions both stop short of actually fixing anything, and trying to build it yourself is not cost-effective.

- LLMs do not know the context of your environment and are “confidently wrong” 60%–70% of the time. Without validation, deploying recommendations can cause major issues.
- Point tools add dashboards and alerts, not the context or authority to act on what they find.
- An internal model only knows your own history, so even after you build it, every new failure mode is still a surprise.

The result: insight piles up faster than anyone can act on it, while cost and performance keep slipping in production.


The Solution

Arvix AI is Unravel's purpose-built optimization engine, built for data platforms from the ground up, not adapted from a general-purpose model.

- Already analyzed 10 billion+ workloads across 100+ enterprises, so it starts with answers, not a training period.
- That scale is exactly what's hardest to reproduce by cloning the idea internally.
- Operates on a continuous loop to find, recommend, test, deploy, and monitor results.

It doesn't stop at telling you what's wrong. It knows your cost structure, performance baselines, and configuration history well enough to fix it.





The Part Other Tools Skip

Most tools stop at Detect or Recommend. Arvix AI is the engine that carries every fix through all six steps below, automatically, from first signal to verified, deployed change.

This is the same engine behind Unravel's FinOps, DataOps, and Data Engineering Agents, each applying it to a different part of your data operations.

01 Ingest

Signals across every platform: billing, query plans, Spark metrics, lineage.

02 Context Graph

Maps every workload, team, and dependency: compute, code, data, platform, business.

03 Match

Checked against a 1B+ workload knowledge base spanning 500+ known failure patterns. The action is already known.

04 Validate

Every fix is tested against real workload behavior before it touches production.

05 Deploy

AutoApply or human-in-the-loop, your choice, per workload or per team, with a full audit trail either way.

06 Monitor

Continuous re-evaluation, with automatic rollback if anything drifts.

Core Capabilities

Query Optimization at Fleet Scale

Rewrites and tunes SQL and code across your entire workload fleet, not just the queries you happen to check.

Granular Cost Attribution

Breaks down spend by workload, team, and job so you know exactly where budget is going, and where it's being wasted.

Compute Rightsizing

Matches cluster and warehouse capacity to actual demand, continuously, not on a quarterly review cycle.

Real-Time SLA Risk Detection

Flags risk before it becomes a breach, with cost and performance guardrails enforced automatically.

The Options Available to the Market Today

Five categories. Each owns a layer. *None closes the loop.*

Every one of these is useful, and most of our customers run two or three of them alongside Unravel. What none of them do is take the finding through rewrite, validation, deployment and rollback without a human doing the work.

CATEGORY	WHAT IT DOES WELL	WHERE IT STOPS
Native platform tools DATABRICKS · SNOWFLAKE · BIGQUERY	Free, built in, and genuinely good at surfacing opportunity inside their own platform.	Recommendations only. Single platform. Under 10% ever implemented.
Data observability MONTE CARLO · ACCELDATA	Tells you when data quality breaks, freshness slips, or a schema drifts.	A different problem. No view of cost or workload performance. Complementary, not competitive.
FinOps and cloud cost CLOUDZERO · APPTIO · FLEXERA	Track, allocate and forecast cloud infrastructure spend across AWS, Azure and GCP.	Cannot see inside the data platforms, which is where most of the overspend is generated.
APM and infra monitoring DATADOG · NEW RELIC · GRAFANA	Monitors CPU, memory, latency and uptime across the estate.	Server level. No Spark stage, micro-partition or slot context. Data workloads are an afterthought.
LLMs and AI assistants CHATGPT · COPILOT · GEMINI	Can suggest a SQL or Spark rewrite when a human pastes in a query.	No environment context. No validation. No deployment. No rollback. Starts from zero every time.

Security & Architecture

Arvix AI runs on job execution and configuration metadata, never your underlying data. Full architecture detail and SOC 2 documentation at trust.unraveldata.com.

About Unravel Data

Unravel Data is an agentic data platform optimization, powered by Arvix AI the engine that continuously optimizes performance, cost, and reliability across Databricks, Snowflake, and BigQuery.

For more information, visit unraveldata.com.

WANT TO SEE THE OPERATOR IN ACTION?

Start with a free Health Check. It runs the same engine against your environment and quantifies the savings sitting in it right now.

Get your results in hours with minimal effort required. Request your free report today at unraveldata.com/health-check-demo/.