acceldata

Data and AI Observability

Trust your data. Trace every AI failure to its source

Most AI Failures Begin with Bad Data

Pipelines break, schemas drift, and numbers stop matching, and for years the job was to catch that before it reached a report. Data observability profiles every source, flag schema and data drift, reconcile across systems, so a problem is caught before anyone downstream inherits it, wherever it lives, on-premises or across clouds. With agents, a quiet data problem becomes a confident wrong action. We added AI observability to watch the agents. When a run fails, it follows back to the data source or pipeline that produced it, in the same console.

No single tool sees your whole estate.

75%

Run four or more data platforms in active use, so no single tool sees the whole estate.

40%

Name governance fragmentation as their biggest cross-platform data challenge.

55%

Would put new budget into data quality, close behind AI and ML.

Source: Independent third-party survey of 40 C-level leaders at Fortune 1000 and Global 2000 firms, April 2026

DATA OBSERVABILITY

Where the data fails

Schema drift
Columns added, dropped, or retyped.

Data drift
Distributions move off their baseline.

Anomaly & range drift
Values fall outside expected bounds.

Freshness & volume drift
Data lands late, partial, or not at all.

AI OBSERVABILITY

Where the AI fails

Prompt drift
A prompt change quietly shifts behavior.

Retrieval / context drift
Wrong or stale context gets pulled in.

Model / output drift
Output shifts after a model or version change.

Accuracy / quality decay
Answers degrade against evaluation sets.

Every Layer Trusted. Traceable to Source.

Your data moves left to right: ingested, transformed, and served, with every stage watched so a problem is caught before anything downstream depends on it. When agents act on that data, the work enters an AI lifecycle. AI observability watches the agents from there, tracing each run step by step, and when one fails, it follows the failure back across the flow to the data that caused it. One connected path, source to output, with data health and agent behavior on the same platform.

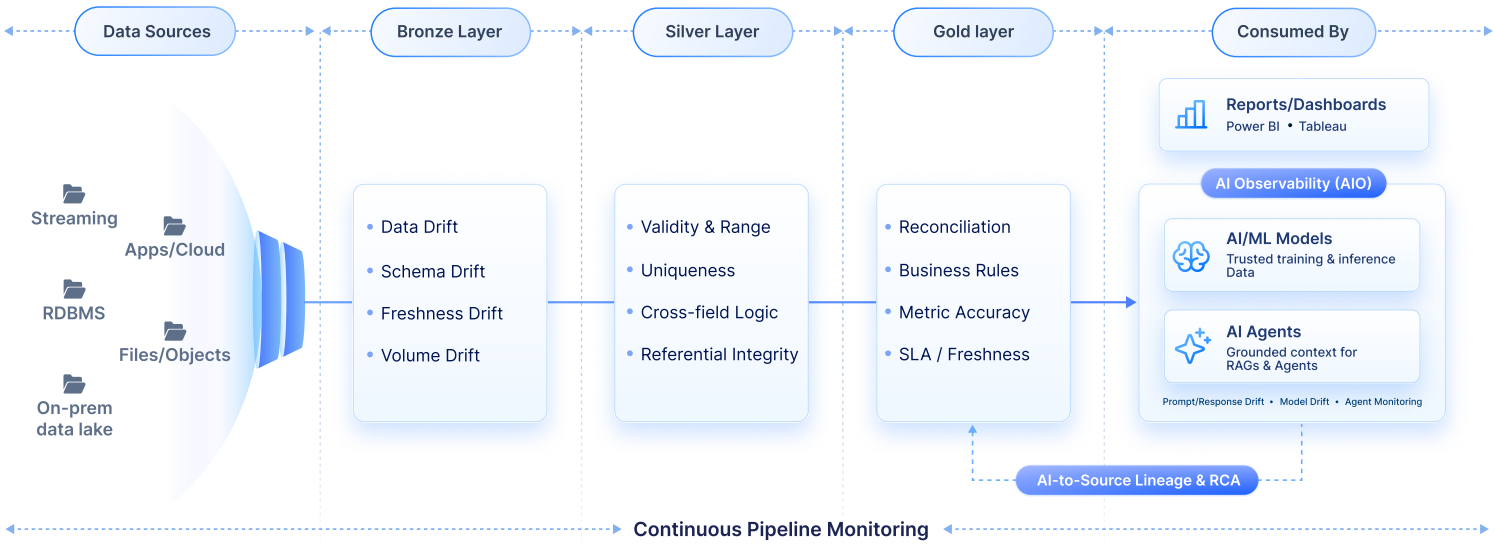

Core Product Capabilities

Data Observability

Profiling and policies
Author by hand, extend with custom code, or auto-recommend with GenAI.

Drift and anomaly detection
Multi-variate detection of drift, schema changes, and freshness violations against baselines.

Reconciliation and governance
Source matched against target, with RBAC and sensitive-column protection.

Lineage to source
Lineage from origin to consumption, back to the asset that caused it.

AI Observability

Full execution tracing
Thread, trace, and span across prompts, model calls, tools, and retrieval.

Online and offline evaluation
Live traffic and regression sets scored for hallucination, faithfulness, and answer relevance.

Prompt management
Versioned prompts and playground, each linked to the traces it produced.

Traceback to the data
One console, so an agent failure follows to the data behind it.

Why Acceldata

Other point solutions tell you an agent failed. We trace it back to the data that caused it, across your whole estate, on one platform.

WHOLE ESTATE COVERAGE

Watches your hybrid estate, not one warehouse. Single-platform monitoring tools stop at their own tables.

SHIFT LEFT DETECTION

Problems surface where data enters, before they reach a report or feed a model. That is where they cost least to fix.

PUSHDOWN AND SCALEOUT

Push checks into the source for zero data movement, or run distributed Spark for heavy cross-source jobs.

> *"Observability goes well beyond the pipeline. It is end to end: the business impact, the SLA impact, the KPIs driving the data. We don't trust something we can't see."*
>
> **VP, Head of Architecture for Data & AI, Fortune 500 enterprise**

> *"It is not the model. It is not the inference that will be successful. It is the data, and the underlying infrastructure."*
>
> **Senior Director, Data and AI, Workday**

Source: Acceldata Autonomous'26 Event

Business Outcomes

Over 1 trillion data interactions daily, in production, for teams that run the same kind of estate you do. Here is what that scale delivers in customer environments.

Faster trust

Models reach production sooner because the data feeding them is trusted before training, not after a failure.

30%

Faster time-to-model

Lower risk

Audit evidence is produced as the estate runs, so governance stops being a fire drill.

Audit-ready

Governance on demand

Less firefighting

Fewer incidents reach production because issues surface at the source, not in the report.

40%

Less pipeline downtime

Visit acceldata.io/demo or sign-up for a 15-day Free Trial