SOLUTION BRIEF

AI observability

See everything your AI does

Cribl

The AI Platform for Telemetry

Introduction

AI observability sounds simple. In practice, it is anything but simple. Most organizations use LLMs to run their business and build apps that integrate with LLMs. Many have large GPU-driven infrastructures to support AI. But very few have visibility into how their AI LLMs and infrastructure truly behave. That gap shows up all at once in cost, quality, security, and governance.

This same lack of visibility also creates new risks. AI telemetry (if it's even being collected) often contains PHI and PII, creating compliance risks. Meanwhile, 59 percent of security leaders confirm or suspect shadow AI, which means many teams only see the AI usage they already know about.

For platform and observability leaders, the takeaway is not "buy another AI dashboard." It is to collect the right data once, shape it well, and make it easy for every each team to explore their data.

The Challenge

Between the exponential growth of AI in your company and the unique nature of how AI operates, ignoring the problem today will snowball into catastrophe tomorrow.

EVERYONE WANTS AI OBSERVABILITY DATA JUST NOT THE SAME DATA.

Platform teams want one collection pass; security wants shadow AI and activity consolidated; FinOps wants cost tied to teams and models; engineering wants to trace failures across the stack, not stitch together three tools. The data's too sensitive and voluminous for one-size-fits-all access.

TRADITIONAL TOOLS WEREN'T BUILT FOR AI.

Models are wired into apps, retrieval systems, copilots, agents, and GPU infrastructure, each emitting telemetry. A model can return HTTP 200 in 200ms and still produce garbage, leak data, or burn money — something uptime dashboards can't see.

ECONOMICS GET UGLY FAST.

Agent traces can carry 200+ KB of detail, and GPUs push telemetry volumes well beyond ordinary infrastructure. Routing it into expensive ingest tools turns observability into a budget problem.

THE HARDEST QUESTIONS ARRIVE LATE.

Security asks what sensitive content hit a model last quarter; FinOps asks who burned budget after the bill lands; compliance wants records the hot tier dropped. Add shadow AI evading controls, and no product solves this.

AI OBSERVABILITY DEMANDS FLEXIBILITY.

OpenTelemetry's GenAI conventions are spreading, but still evolving — schema churn comes with the territory. Your approach must evolve.. Your telemetry will keep changing, and your management approach needs to change with it.

Cribl: The AI Platform for Telemetry

Cribl gives organizations the choice, control, and flexibility to collect, route, enrich, and store telemetry from any source to any destination, at the scale and economics modern AI workloads demand.

A telemetry-first approach to AI observability

Cribl treats AI observability as a telemetry problem, not another monitoring silo. Rather than each team collecting its own copy of AI telemetry — or funneling every signal into one destination — Cribl provides:

- A control layer for collecting, shaping, governing, and routing telemetry
- A place to analyze the full picture across tools and storage tiers
- A platform to build and use apps that help you answer the questions you have

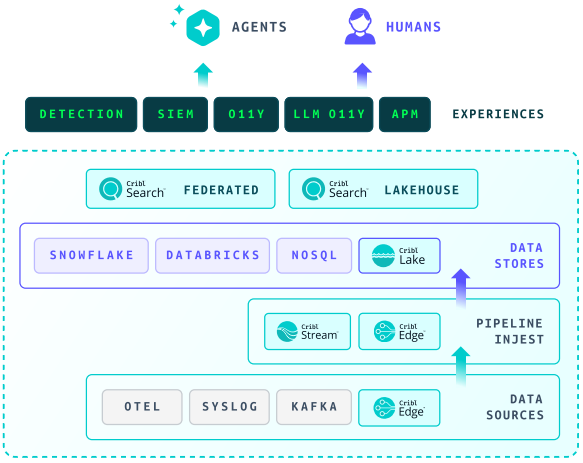

Shared Telemetry Infrastructure

Full-fidelity history, on demand

Effective AI observability depends on full-fidelity history and strong analysis. Cribl Lake stores LLM and infrastructure telemetry at low cost; Cribl Search turns it into insight via ad hoc investigation, dashboards, and scheduled searches — so teams can ask long-tail questions months later without rehydrating to a hot tier.

Trust and safety, built in

AI telemetry needs careful handling before it crosses trust boundaries. Cribl Guard masks prompts, completions, API keys, PII, and PHI in flight, while preserving fields needed for debugging, governance, and cost monitoring. Cribl can also correlate instrumented usage with network egress to surface ungoverned AI use.

One telemetry layer, every team

- SRE — performance and reliability context
- Security — risky usage and policy-relevant events
- FinOps — usable chargeback inputs
- AI teams — visibility into model and infrastructure performance

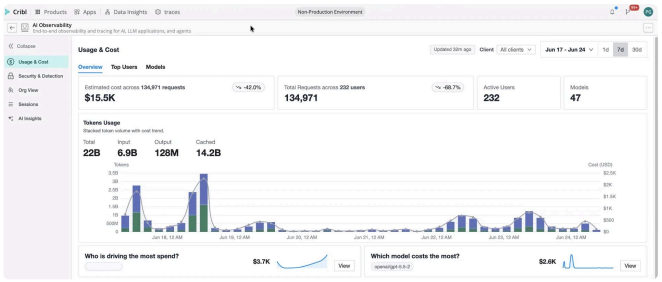

Summary

AI observability is the discipline of turning AI telemetry into something every team can trust, afford, and use. As AI adoption spreads across applications, agents, and GPU infrastructure, Cribl give your teams the choice, flexibility and control to collect AI observability telemetry once, govern it well, route it wisely, and give people clear visibility and answer to the questions they have about AI observability.

Cribl® | The **AI Platform** for Telemetry

Learn more at cribl.io | Try Cribl Sandboxes | Join our Slack community | Follow us on LinkedIn and X