

Total Cost of Ownership for Operationalizing Agents

External LLMs vs. Fiddler Trust Models at enterprise scale

Every observability platform promises visibility into your agents. But most don't tell you the full cost of that visibility, including a hidden cost called the Trust Tax.

How Most Platforms Evaluate Agents

When an agent generates a trace, it needs to be scored for quality, safety, and performance. The most common approach is to send each trace to an external LLM (OpenAI, Anthropic, etc.) to act as a judge. This means every trace scored is an API call to a third-party model provider, and that cost shows up on your bill, not your tooling vendor's.

This approach is called LLM-as-a-Judge. When the judge is an external LLM, the dependency on third-party API calls is what drives up your total cost of ownership (TCO).

The Cost of Evaluating with External LLMs

- 1 Your trace is generated by your agent
- 2 API call to external LLM provider for scoring (OpenAI, Anthropic, etc.)
- \$ You pay the Trust Tax

The Hidden Costs of Calling External LLMs

- Risk Gaps**

To control costs, some teams may consider sampling, but the traces you skip are often the ones that matter most: jailbreak attempts, policy violations, hallucinations, edge-case failures. These are low-frequency, high-impact events that sampling could potentially miss.
- Operational Overhead**

The engineering effort to set up, manage, and maintain your evaluation infrastructure, whether that's orchestrating external API calls or standing up your own models. Your team carries the burden of prompt versioning, scoring calibration, model hosting, and ongoing maintenance.
- The Trust Tax**

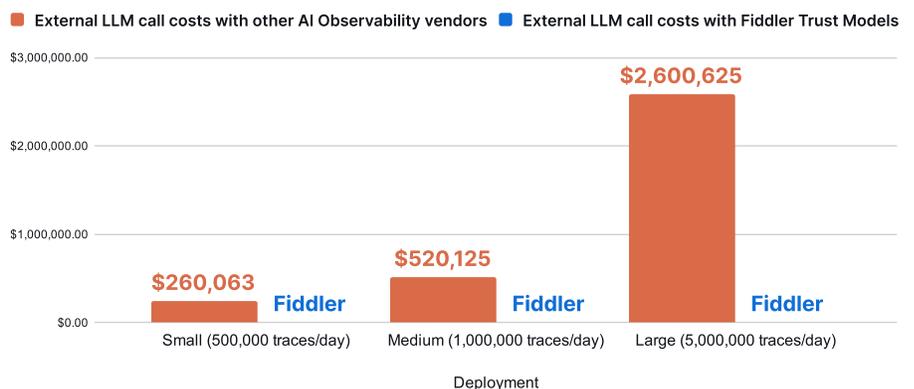
You are charged every time a trace is scored via an external LLM API call. This shows up on your invoice. At enterprise scale, it compounds fast.

What the Trust Tax Looks Like at Scale

These are potential external LLM API costs you pay annually, in addition to your tooling fees.

* Calculations based on Open AI GPT 5 mini.
1 trace = 1 API call to GPT 5 mini.

Contact sales@fiddler.ai to receive a custom calculation.



Fiddler Trust Models: Evaluate and Observe Agents Without the Trust Tax

Fiddler Trust Models are specialized, task-specific models built in-house and deployed securely in your environment. They score agent and LLM prompts and responses at runtime for hallucination, toxicity, jailbreaks, PII/PHI exposure, and other critical risks.

Trust Models are built to cover a range of use cases:

Out-of-the Box Models

- Hallucination detection, safety scoring, toxicity, jailbreak detection, and PII/PHI identification
- Ultra low-latency and task-specific

Customizable Models

- Enterprises submit prompts to create domain-specific evaluators
- Fully managed, handles 300K+ daily events without the burden of infrastructure management

Why Sampling Doesn't Solve the Trust Tax

Sampling reduces your bill but introduces risk gaps. **Fiddler Trust Models cover 100% of traces by removing the cost barrier that draws teams to sample in the first place.**

Fiddler Trust Models Power AI Observability and Security



Evaluation

Test and benchmark agents before they go live. Run evals against test sets, compare model versions, and validate guardrail thresholds before launch.



Observability

Continuously score and monitor every trace in production. Surface issues in real time, diagnose root causes, and trigger alerts.



Guardrails

Enforce safety policies in real time across input, execution, and output, preventing violations before they occur.



Analytics

Roll up aggregate reports down to granular insights across agents for a single-pane-of-glass view of behavior, risk, and performance.

Trusted by Industry Leaders and Developers

"Fiddler delivered unified observability, protection, and governance across agents and predictive models, making it fundamental to our AI strategy."

Karthik Rao, CEO, Nielsen

 Nielsen

 Thumbtack

 AIG



 LENDINGPOINT



CSAA Insurance Group,
a AAA Insurer

 Elevance
Health

 mastercard.

Fiddler is the all-in-one AI Observability and Security platform for responsible AI. Our evaluations, monitoring and analytics capabilities provide visibility, context and control across development and production. This gives teams actionable insights to build better, more reliable AI agents, and GenAI and ML applications. An integral part of the platform, the Fiddler Trust Service provides quality and moderation controls for AI agents and GenAI applications. Powered by cost-effective, task-specific, and scalable Fiddler-developed Trust Models — they deliver the fastest guardrails in the industry. Fiddler offers flexibility in secure deployment options through cloud and VPC environments.

Fortune 500 organizations use Fiddler to scale AI agents, GenAI, and ML deployments. This helps them deliver high performance AI, avoid costly AI risks, and maximize ROI.