

Build Responsible AI with AI Observability and Security

Fiddler partners with you in creating a long-term framework for responsible AI.

Fiddler is a pioneer in AI Observability and Security — the foundation which gives enterprises the confidence to ship more predictive, generative, and agentic applications and models into production safely and responsibly. We help instill trust and transparency in your agentic, LLM, and ML applications — factors that directly impact your return on investment (ROI). By driving AI governance, risk management, and compliance, Fiddler supports the advancement of responsible AI practices. AI Observability is reliant not only on metrics, but also on how well issues can be explained when something eventually goes wrong.

Fiddler empowers your Data Science, Engineering, Trust & Safety, and Security teams to monitor, analyze, and protect agentic, LLM, and ML applications.



Deliver High Performance AI



Reduce Costs and Increase ROI



Be Responsible with Governance



IAS

Thumbtack



Build Trust into AI with Fiddler

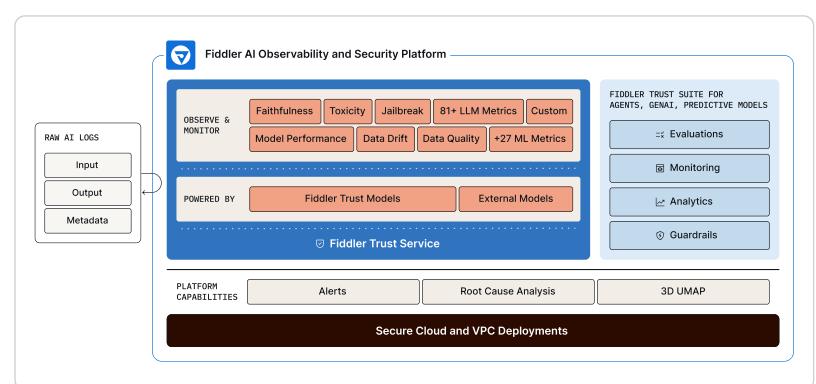
All-In-One AI Observability and Security Platform:

A unified platform that brings together agentic, LLM, and ML observability, Fiddler Trust Service, Guardrails, and contextual model analytics.

Actionable Diagnostics: Deep root cause analysis and actionable application and model insights for quick issue resolution and model improvement.

Built for the Enterprise and Government:

Enterprise-grade scalability and stability to support secure agentic, LLM, and ML deployments in SaaS and VPC environments, including AWS GovCloud for the government.



"One of the things that was appealing to IAS about Fiddler was its ability to customize the monitoring to specific model type, data volume and desired insights. Additionally, the dashboard views, automated alerting and ability to generate audit evidence also factored into the decision to work with Fiddler."

Kevin Alvero, Chief Compliance Officer, Integral Ad Science

Key Capabilities

 Monitoring	 Analytics	 Security
<p>Monitor predictive models and LLMs in production and manage performance, correctness, and safety metrics at scale in a single pane of glass. From alerts to root cause analysis, pinpoint areas of application and model underperformance and minimize business impact. You can also find quick answers to the root cause and the “why” behind all issues.</p> <p>Plug Fiddler into your existing agentic, LLM and ML tech stacks for consolidated monitoring to:</p> <ul style="list-style-type: none">• Gain operational efficiencies through faster time-to-market at scale.• Reduce costs by improving MTI and MTTR for AI applications.• Foster responsible innovation by safeguarding agents, LLM and ML applications and data.	<p>Analytics must deliver increased transparency and actionable insights that power data-driven decisions. To improve predictions, market context and business alignment must be baked into modeling so results reflect the needs and challenges of your business.</p> <p>Use descriptive and prescriptive analytics from LLMs and ML models to make decisions so you can:</p> <ul style="list-style-type: none">• Deploy higher ROI models to increase revenue.• Align decisions to stay in lockstep with business needs.• Respond quickly and improve applications and models when market dynamics shift.	<p>Safeguard AI applications with enterprise-grade security and low-latency guardrails. Fiddler enables secure deployment in SaaS and VPC environments, preventing data leaks, prompt injection attacks, and jailbreak attempts while upholding trust and safety.</p> <p>Proactively protect against security risks:</p> <ul style="list-style-type: none">• Prevent security threats like prompt injection and adversarial attacks.• Reduce risk exposure with compliance-driven monitoring and alerts.• Ensure data privacy with secure deployments in SaaS and VPC environments.

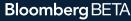
Agentic and LLM Use Cases

AI Chatbot Monitoring Content Summarization
LLM Cost Management Internal Copilot Monitoring
Multi-Agent Interdependency Impact
Agentic Hierarchy Monitoring

ML Use Cases

Anti-money Laundering Fraud Detection
Credit Scoring Investment
Decision-Making Underwriting
Churn Detection

AI Governance Risk Management Compliance Security

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.