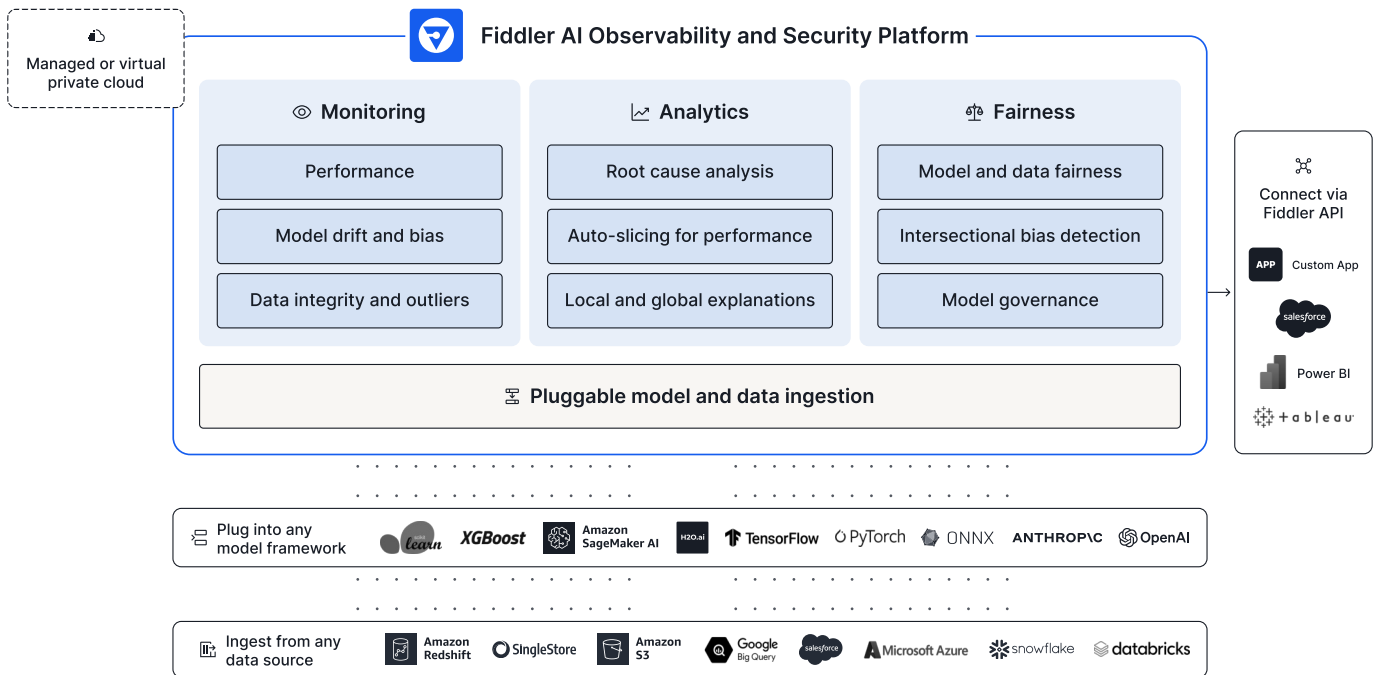


Build Responsible AI with Mission-Critical AI Observability

Fiddler partners with you in strengthening National Security

Fiddler is a pioneer in AI Observability and Security — the secure, deployment-ready foundation you need to standardize LLM and MLOps practices. Fiddler helps government agencies ensure AI is mission-ready with advanced monitoring and security, enabling rapid model adaptation to evolving operational and national security needs — while embedding transparency, security, and trust into every AI application. AI Observability is reliant not only on metrics, but also on how well issues can be explained when something eventually goes wrong.



With contributions from Fiddler AI, The U.S. Navy slashed model retraining and update time by 97%, from 6 months to just a few days. To achieve this, Fiddler helped established a robust MLOps pipeline, automating issue detection like model drift and eliminating time-consuming manual work for developers.

Government Agencies Trust Fiddler with AI Innovation



Government Trusted: Trusted by government agencies that aim to accelerate AI innovation to strengthen national security.



Air-gapped Deployment: Air-gapped deployments to accelerate mission-critical projects that require little or no connectivity.



Secure Deployment: Available on AWS GovCloud, one of the largest secure cloud solutions that addresses the most stringent U.S. government security and compliance requirements.



Edge Support: AI Observability for edge devices, giving agencies widespread flexibility for high performance monitoring and explainable AI.



Petabyte-scale: Enables government agencies with petabyte scalability to support large AI deployments.



Trusted Advisors and Partners: White glove support from Fiddler's AI science experts, dedicated to ensuring the success of your mission-critical LLM and ML projects.

Key Capabilities



Monitoring

Monitor predictive models, NLP, CV, and LLMs in pre and post-production, and manage all performance metrics at scale in a unified dashboard. From alerts to root cause analysis, pinpoint areas of model underperformance and minimize impact. You can also find quick answers to the root cause and the "why" behind all issues.

Plug Fiddler into your existing ML and LLM tech stacks for consolidated monitoring to:

- Receive real-time alerts on potential threats from image data inconsistencies.
- Improve target precision with high performance model monitoring across different battlefield environments.
- Monitor changes in traffic, signals, and communication patterns.



Analytics

Analytics must deliver actionable insights that power data-driven decisions. To improve predictions, market context and business alignment must be baked into modeling so results reflect your needs and challenges.

Fiddler's proprietary Explainable AI (XAI) technology enhances analytics by providing complete context and visibility into ML model outputs, from training to production.

Implementing descriptive and prescriptive analytics from ML models and LLMs enables you to:

- Enhance decision-making and adaptability by integrating human-in-the-loop processes.
- Improve model performance and threat detection by identifying anomalies through advanced analytics and explainability.



Security

Fiddler Guardrails is the fastest solution on the market, proactively moderating risky LLM prompts and responses before they cause harm. Leveraging the scoring of the Fiddler Trust Service, guardrails deliver fast, cost effective, and secure moderation of harmful risks such as hallucinations, toxicity, prompt injection attacks, and more.


- With a <100ms latency, guardrails ensures enterprises can quickly detect and moderate LLM issues.
- Fiddler can be deployed in VPC or air-gapped environments, maintaining compliance and protecting sensitive data.
- Fiddler Guardrails is efficient, with an average computational overhead 6x lower than closed source foundational models*.

*Fiddler Trust Models are benchmarked against publicly available datasets

Responsible AI Practices for Government Agencies


Responsible AI is the practice of building transparent, accountable, ethical, and reliable AI. The first step is detection and mitigation of bias in tabular and unstructured datasets and ML models, but you must also support internal governance processes and reduce risk through human involvement.

Build and deploy responsible AI solutions with bias detection and fairness assessment in order to:




Reduce Risk

Instill trust with continuous AI monitoring and human decision-making with ML.



Support Internal Oversight Teams

Provide visibility and AI governance.



Mitigate Bias

Through the detection, comparison, and measurement of fairness metrics.

Use Cases

- Chatbot Monitoring and Guardrails
- Image Explainability for Autonomous Vehicles
- Model Monitoring Sensor Data
- Cyber Security ML Monitoring
- NLP and CV Monitoring for Intelligence and Surveillance



INSIGHT PARTNERS Lightspeed LUT CISCO investments Capgemini DALLAS VENTURE CAPITAL ALPHA SQUARE SAMSUNG NEXT ESSENCE
Dentsu Ventures iqt LOCKHEED MARTIN Bloomberg BETA Haystack alexa fund mozilla ventures alteryx

Fiddler is the all-in-one AI Observability and Security platform for responsible AI. Monitoring and analytics capabilities provide a common language, centralized controls, and actionable insights to operationalize production ML models, GenAI, and LLM applications with trust. An integral part of the platform, the Fiddler Trust Service provides quality and moderation controls for LLM applications. Powered by cost-effective, task-specific, and scalable Fiddler-developed trust models — including air-gapped deployment for secure environments — it delivers the fastest guardrails in the industry.

Fortune 500 organizations use Fiddler to scale LLM and ML deployments to deliver high performance AI, reduce costs, and be responsible in governance.