Data Sheet

fiddler

# 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, providing a secure, deployment-ready foundation to gain visibility, context, and control over agentic, generative, and predictive AI across development and production. Fiddler helps government agencies ensure AI is mission-ready with advanced evaluation, monitoring, analytics, and security. This enables rapid model adaptation to evolving operational and national security needs while embedding transparency, trust, and accountability 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.

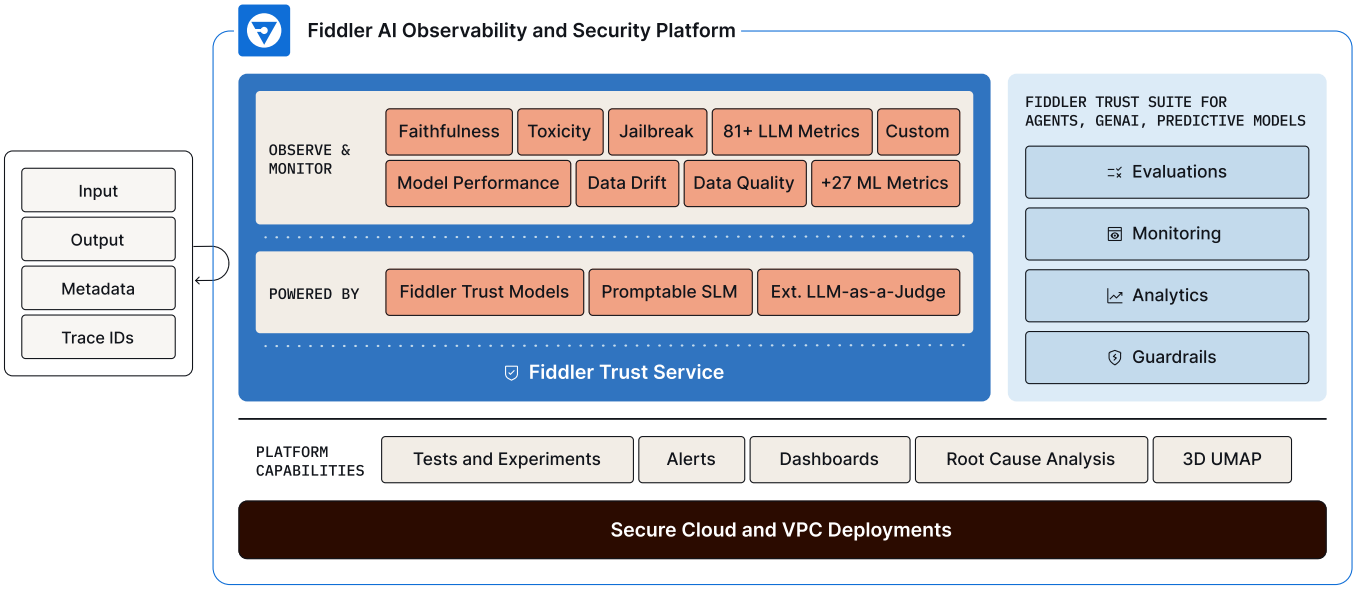


iqt IN-Q-TEL

DEFENSE INNOVATION UNIT

LOCKHEED MARTIN

NIWC PACIFIC

With contributions from Fiddler AI, The U.S. Navy slashed model retraining and update time by 97%. 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 agentic, GenAI and ML projects.

## Key Capabilities

### Evaluations

Evaluate agents and LLMs before deployment against mission scenarios to assess accuracy, performance, and behavior, and improve systems before launch. Build golden and challenger datasets, run out-of-the-box or custom evaluators, and compare prompts and responses side by side to determine mission deployment readiness.

With Fiddler Evaluations, you can:

- Test different scenarios: standard and adversarial cases to improve performance and behavior.
- Run experiments with different prompts and responses and compare side-by-side to reach better outcomes.
- Stress-test edge-cases to surface weaknesses and points of failure modes early.

### Monitoring

Monitor agents, LLMs, and predictive models (tabular, NLP, and CV) 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 agentic, LLM, and 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, from granular details to aggregated views, that power data-driven decisions and autonomous action. To improve predictions, deliver accurate and safe outputs, and enable intelligent automation, organizations need comprehensive analytics that roll up detailed metrics into complete visibility, meaningful context, and actionable control.

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 agents, LLMs, and ML models 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*.

## Responsible AI Practices for Government Agencies

Responsible AI is the practice of building transparent, accountable, ethical, and reliable AI, from traditional machine learning models to generative AI and autonomous agents. Begin with detecting and mitigating bias in datasets and ML models, and ensuring generative AI outputs are correct and safe, while establishing robust governance practices and maintaining human oversight.

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 models, generative AI, and agents.

### Support Internal Oversight Teams

Provide visibility and AI governance.

### Mitigate Bias and Risks

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

## Use Cases

- Chatbot Monitoring and Guardrails
- Image Explainability for Autonomous Vehicles
- NLP and CV Monitoring for Intelligence and Surveillance
- Model Monitoring Sensor Data
- Cyber Security ML Monitoring
- Agentic Systems with Embedded LLM and ML Models

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.

fiddler.ai sales@fiddler.ai