

Agentic, Generative AI and ML Observability and Security for Building Responsible AI

Fiddler serves as the secure, deployment-ready foundation to standardize agentic, LLM and MLOps practices, empowering government agencies to ensure AI is mission-ready. With advanced monitoring and security, agencies can quickly adapt models to meet evolving operational demands—all while embedding transparency, trust, and accountability into every AI application.

Observe and Secure AI Applications

-  **Evaluate:** Test and experiment agents and GenAI applications before deployment against mission scenarios to assess accuracy, performance, and behavior, and improve systems before launch.
-  **Monitor:** Comprehensive model monitoring with customizable alerts, dashboards, and root cause analysis.
-  **Analyze:** Deep root cause analysis and explainability provide complete context and visibility into inputs, outputs, and decision-making across the entire AI lifecycle, from training and testing to production.
-  **Protect:** Fiddler Guardrails moderate LLM risks with industry-leading <100ms latency,* with support for government deployment options to ensure data remains secure.

Technical Competencies

- Agentic Observability
- Agentic Evaluation
- LLM Application Monitoring
- LLM Application Guardrails
- ML Model Monitoring (Tabular, Computer Vision, Natural Language)
- ML Model Output Explainability

Company Snapshot

UEI Code **NEKQEH58N77**

CAGE Code **8H2E0**

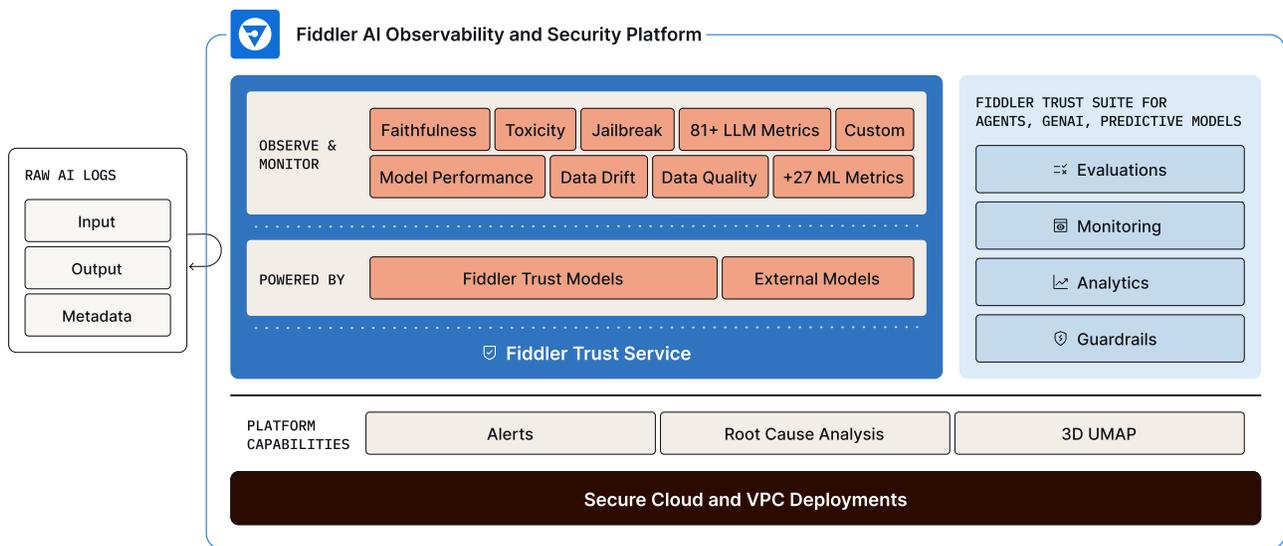
NAICS Codes

541511	Custom Computer Programming Services (Primary)
541519	Other Computer Related Services
541512	Computer Systems Design Services
541715	Research and Development in the physical, Engineering and Life Sciences
513210	Software Publishers

Gov. Business Contact: Amit Paka
Title: Chief Operating Officer
Email: amit@fiddler.ai

Fiddler Labs Inc.,
550 California Ave, Ste 200
Palo Alto, CA 94306, USA





Differentiators



Unified Solution for Agents, GenAI and Predictive ML

Fiddler offers a single pane of glass for monitoring and securing agents, GenAI, and predictive ML applications.



Fiddler Trust Service

A series of purpose-built, “batteries-included” Fiddler Trust Models that enable high quality, low latency LLM scoring in live environments.



Fastest Guardrails in the Industry

Moderate hallucinations, toxicity, and jailbreak attempts with <100ms latency.*



Cost Effective

Fiddler Trust Models reduce overhead by 18x compared to closed source foundational models.



Secure

Deploy Fiddler in your secure VPC on AWS GovCloud.



Case Study: U.S. Navy

The U.S. Navy Expeditionary Missions Office leverages Automatic Target Recognition (ATR) machine learning models to identify potential threats, but faced challenges monitoring and improving these models post-deployment.

To address this, the Navy and Defense Innovation Unit launched Project AMMO, selecting Fiddler to help build an MLOps pipeline — a system for monitoring model performance over time and automatically surfacing potential issues and areas for improvement.

This helped the Navy proactively identify and mitigate model drift, identify anomalies, and enhance human decision-making in post-mission analysis.

The partnership with Fiddler helped lead to:

- **97% reduction in model update time** — from 6 months to 3 days.
- **Improved model explainability** to help human decision makers.
- **A transition to production** with NIWC after a successful prototype.

* Dependent on input size, geographical location, system load, or other infrastructure variability.

