

The Fiddler Control Plane for AI Agents


Continuous oversight and real-time enforcement, from the creation layer to production.

Agents are first-class enterprise actors, making autonomous decisions across your operations: reading files, calling tools, and pulling context faster than any reviewer can keep up with. However, most organizations don't fully know what they're accessing, spending, or exposing.

Fiddler is the system of trust for your agents. The Fiddler AI Control Plane gives you visibility, context, and control across the AI lifecycle, covering first-party, third-party, and coding agents.


Visibility. Context. Control.

The Fiddler Control Plane delivers three enterprise benefits, deploy high-performance agents, prevent compliance and security risks, and maximize ROI, through five core capabilities.



The Only Inline Enforcement at the Agent's Request and Response Path

For coding agents, Fiddler integrates with the LLM gateway you already operate. No new gateway, no agent-side integration.



Integrate Fiddler through the existing LLM gateway in your AI stack:

LiteLLM

AgentGateway

Kong Gateway

One Control Plane. Through the Gateway You Already Run



Adoption, Performance and Cost Observability

Agentic fleet intelligence across every agent, developer, token, and dollar.

- Native OpenTelemetry, no instrumentation; full telemetry through your gateway
- From individual developer to org-wide rollout
- Cost, performance, and security in a single pane
- Contextual insights for executives across every business unit



Bi-Directional PII/PHI Detection and Redaction

Detects and redacts sensitive data in both directions, whether prompted by a developer or returned by a tool.

- Enforces on both request and response; traditional tools catch only the request
- Model-based detection via Centor Models, not pattern matching
- Zero data egress: runs in your environment, no Evaluation Trust Tax



Inline Secrets Detection and Redaction

Catches what commit-time scanning misses: credentials read from .env files, configs, or MCP endpoints.

- Detects leaks in developer prompts, file reads, and tool returns
- Stops leaks before data moves anywhere
- Runs with bidirectional detection at the same gateway point

Flagging is Not Enforcing

Vendors in this space flag risky behavior with evaluators and dashboards. A score logged after the agent acts documents the incident; it doesn't prevent it. Fiddler observes and enforces on the same request: every interaction captured for visibility and gated before any data moves.

Trusted by Industry Leaders and Developers



CSAA Insurance Group,
a AAA Insurer



Fiddler is the AI Control Plane, the system of trust, for first-party and third-party agents from the creation layer to production. Continuous evaluation, reliable monitoring, enforceable policy, and auditable governance give enterprises the centralized controls and actionable insights to scale AI with trust. Integral to the platform are the secure Fiddler Centor Models, which power the industry's fastest guardrails and evaluations that require low latency and cost-effectiveness or complex reasoning with accuracy. Compared to external LLMs, Fiddler Centor Models offer a low TCO with no hidden costs.

Fortune 100 organizations use Fiddler to deliver high performance agentic and predictive applications, protect from costly risks, and maximize ROI.