

NEW ENTRANTS

Agent Observability and Runtime Enforcement Take Center Stage

New entrants, week of 2026-09-29 to 2026-10-05.

Scope: 2026-09-29 to 2026-10-05

Generated 2026-10-05

Report 6dac29a8-e662-4a3e-a849-cfd331091717

10
NEW ENTRANTS

6
DEEP-DIVED

8
TRACKED-PRODUCT MOVES

Week of 2026-09-29 to 2026-10-05.

NEW THIS WEEK

Ten tools entered the pack this week, with 5 landing in evaluation and observability, a category that is clearly seeing competitive heat. The cluster signals a market converging on agent-native monitoring as a distinct discipline, separate from classical ML observability.

[AgentOps](#) leads with a freemium, SaaS-plus-on-prem posture and a compliance story that stretches to SOC 2, HIPAA, and SSO, unusually complete for an early-stage startup. Its "time travel debugging" and session replay are differentiators for teams debugging non-deterministic agent behavior. [Fiddler AI](#) arrives with inline guardrails rated under 80ms alongside PII and PHI redaction, positioning it as a single-plane answer to both observability and runtime protection. [Braintrust](#) covers the broadest compliance surface of the cohort: SOC 2, HIPAA, GDPR, SSO, and data residency, making it immediately relevant to regulated-sector buyers. [Arize Phoenix](#) and [Arize AI](#) enter together as sibling products: Phoenix is the open-source debugger; Arize AI the commercial platform adding online evals, root-cause analysis, and an Alyx AI agent for debugging workflows. Arize AI carries SOC 2, ISO 27001, HIPAA, GDPR, and PCI, the strongest compliance stack in the cohort.

On the security side, [Doberman](#) takes a fail-closed, policy-as-code stance at MCP proxy level, and [NVIDIA's Open Agent Safety Platform](#) introduces hardware-attested telemetry and agent quarantine in milliseconds. Both signal that runtime enforcement, not just logging, is becoming a buyer requirement. The two enterprise knowledge entrants, [Movo](#) and [qKnow](#), and the analytics agent [nao](#), round out the cohort at a solid fit rating.

MOVES ON TRACKED PRODUCTS

The week's 8 tracked-product moves are dominated by [Open Agent Safety Platform](#), which collected three separate press items spanning TechCrunch, NVIDIA's own developer blog, and an eesel.ai explainer, all between 2026-09-29 and 2026-10-02. The TechCrunch piece ([link](#)) flags an industry-coalition angle with a notable absence: OpenAI. For a regulated buyer evaluating vendor ecosystem stability, that gap is worth tracking before committing to the platform's governance model.

[Braintrust](#) had an active week on the compliance front, publishing both a SOC 2 Type II announcement ([link](#)) and a security and data control deployment model post ([link](#)) on the same day it entered the pack. The timing is deliberate positioning for enterprise procurement cycles.

[Langfuse](#) appears in an InfoWorld piece on new SLIs for AI-native systems, reinforcing its relevance to teams building observability frameworks beyond conventional latency and error-rate metrics. Hoop and Nuwax Agent OS each registered a GitHub activity move, thin signals that suggest development velocity rather than any strategic announcement.

WATCH

The NVIDIA [Open Agent Safety Platform](#) coalition story warrants close attention: three press hits in under a week, an explicit gap with OpenAI noted by TechCrunch, and no pricing or procurement path yet visible in the pack. Clarity on coalition membership and commercial availability will matter for any enterprise piloting agent infrastructure.

[Braintrust](#) is worth a deeper procurement review. It entered the pack and announced SOC 2 Type II on the same day, which suggests it was timing the pack entry to coincide with compliance readiness. Its compliance breadth, SOC 2, HIPAA, GDPR, SSO, and data residency, is the strongest argument for a regulated-sector shortlist, but pricing is currently listed as unknown, a gap that needs to close before a serious evaluation can begin.

Bottom line: This week's cohort establishes agent-native observability and runtime enforcement as the two fastest-moving fronts in the AI tooling market, with 5 of 10 entrants landing in evaluation and observability and two dedicated security tools signaling that logging alone no longer satisfies buyer requirements. Compliance depth varies sharply: Arize AI and Braintrust carry the strongest stacks, while NVIDIA's coalition platform and several open-source entrants remain thin on procurement-ready signals. The single most important next step is resolving Braintrust's unknown pricing, as its compliance breadth makes it the most immediately shortlist-eligible tool in the cohort for regulated buyers once that gap closes.

PRODUCTS

PRODUCT	MATURITY	DEPLOYMENT	PRICING	FEATURES
AgentOps AgentOps · An open-source observability framework for AI agents that tracks usage ...	Early stage startup	saas, on_prem	freemium	agent tracing, time travel debugging, session replay, token tracking, cost tracking
Doberman DobermanCore · An AI guard dog that sits at runtime to gate inputs, outputs, and tool ...	Open source project	open_source, on_prem, free desktop		runtime tool call gating, pass/auth/block verdicts, mcp proxy, host hooks, fail-closed enf ...
Fiddler AI Fiddler Labs · An AI observability and governance platform for monitoring model perfo ...	Growth stage startup	saas, vpc, on_prem	freemium	agent monitoring and observability, continuous evaluations, inline guardrails under 80ms, ...
Open Agent Safety Platform NVIDIA · A security platform and reference system design to set safer boundarie ...	Big tech	open_source, on_prem	free	secure runtime boundary, agent policy enforcement, out-of-band watchdog, agent quarantine ...
Braintrust Braintrust · An end-to-end platform for evaluating and improving LLM applications w ...	Growth stage startup	saas, vpc	unknown	agent trace observability, llm and human evals, production log search, pattern discovery, ...
Arize Phoenix Arize AI · An open-source observability toolkit from Arize for debugging and anal ...	Open source project	open_source, saas, on_prem	freemium	agent tracing, llm evaluation, llm-as-judge annotations, human review labeling, experiment ...
Arize AI Arize AI · An AI observability and evaluation platform for monitoring and improvi ...	Open source project	saas, open_source, api	unknown	agent tracing, online evaluations, span and session evals, prompt testing, root cause anal ...

Movo himovo · A self-hosted enterprise Agent platform built on DeepSeek Harness, fea ...	Open source project	on_prem, open_source, free desktop	enterprise knowledge search, multi-round research with citations, document parsing pdf doc ...
nao getnao · An open-source analytics agent that creates context via CLI and deploy ...	Open source project	open_source, on_prem free	natural language to insights, chat data visualization, agent performance unit testing, con ...
qKnow qiantongtech · An open-source agent platform for enterprise knowledge intelligence pr ...	Open source project	open_source, on_prem freemium	knowledge graphs, rag question answering, bot building, visual workflow orchestration, ent ...

MOVES ON TRACKED PRODUCTS

- Braintrust: news "Braintrust achieves SOC 2 Type II compliance" on 2026-10-05
- Braintrust: news "Braintrust publishes security and data control deployment model blog post" on 2026-10-05
- Hoop: news "hoophq/hoop" on 2026-10-05
- Langfuse: news "Observability for AI-native systems: New SLIs beyond latency and error rate | InfoWorld" on 2026-10-05
- Nuwax Agent OS: news "nuwax-ai/nuwax" on 2026-10-05
- Open Agent Safety Platform: news "NVIDIA Open Agent Safety Platform: A Reference for Continuous In-Silicon Agent Monitoring (NVIDIA De ..." on 2026-10-02
- Open Agent Safety Platform: news "NVIDIA Open Agent Safety Platform explained: OpenShell, Sentry, cost (eesel.ai)" on 2026-09-30
- Open Agent Safety Platform: news "Here's why OpenAI is absent from Nvidia's industry-wide effort to end rogue AI agents (TechCrunch)" on 2026-09-29

Methodology: the data half of this report (counts, directions, spans, coverage statements) is computed deterministically from the AI Atlas database. The narrative half is drafted by an AI model (Claude) over that frozen evidence pack only, then passed through a citation gate: any link the pack cannot vouch for is removed before rendering. Statistics in the prose come from the computed figures. Generated 2026-10-05.

The AI Atlas is an orientation tool, not a research product: it maps where the disagreement is and what evidence would move it. Nothing here is investment, legal, or professional advice.