

NEW ENTRANTS

Agentic Stack Strengthens Across Five Regulated-Buyer Categories

New entrants, week of 2026-09-22 to 2026-09-28.

Scope: 2026-09-22 to 2026-09-28

Generated 2026-09-28

Report 7185a203-ec31-40da-85a6-f6bee52e5bad

8

NEW ENTRANTS

5

DEEP-DIVED

12TRACKED-PRODUCT MOVES

Week of 2026-09-22 to 2026-09-28.

NEW THIS WEEK

Eight tools entered the pack this week across five categories: 2 coding assistants and agents, 2 agent platforms and orchestration, 2 enterprise search and knowledge assistants, 1 guardrails and security, and 1 model gateways and LLMOps.

The most strategically distinctive entrant for a regulated buyer is [T1](#), a Claude Code variant from Ginlix AI purpose-built for financial market analysis. Its feature set is unusually domain-specific: SEC filing analysis, options chains, treasury yield curve, macro data, and earnings transcripts sit alongside parallel research agents and news sentiment. Combined with `no_training_on_customer_data` and `data_residency` compliance flags, it is immediately credible for capital markets teams with data-handling obligations.

[T3](#) addresses a gap that most agentic stacks leave open: runtime enforcement. Its no-code sidecar proxy model means deployment does not require application code changes, and its feature set, covering PII masking, destructive command blocking, human approval routing, and session replay logs, maps directly to controls an internal audit function would ask for. SOC 2 and audit logs are confirmed.

[T5](#) and [T6](#) both enter the enterprise search category as free, on-prem-deployable open-source platforms. PipesHub leads on connector breadth (50+) and MCP support; Nuwax Agent OS differentiates on agent sandbox deployment and browser automation. [T7](#) brings the strongest compliance posture of the week: HIPAA, GDPR, SOC 2, and data residency, relevant if document pipelines touch regulated content.

[T4](#) and [T2](#) round out the orchestration layer, with Plano adding jailbreak detection and centralized policy enforcement that LangChain does not surface.

MOVES ON TRACKED PRODUCTS

Twelve moves were recorded on tracked products this week, with [T5](#) generating the most signal: three separate items confirmed its Apache 2.0 GitHub publication ([source](#)), its self-positioning as "[The Open Source Glean Alternative](#)", and its entry into the AI2 Incubator program ([LinkedIn](#)). For a buyer evaluating Glean, this is a direct, cost-model alternative worth a structured comparison.

[T1](#) received an external listing on aikii.org ([source](#)) characterizing it as an open-source financial Agent OS, consistent with the pack's own description.

[T4](#) had a community feature request open on per-call agent cost settlement at the gateway layer ([issue](#)). This is an early signal of demand for chargeback-style cost attribution across agent calls, a capability gap that matters for enterprise procurement and internal cost allocation.

Encord saw two moves: a [pricing update](#) formalizing VPC and on-prem as add-on tiers, and a [security page](#) confirming SOC 2 completion and HIPAA compliance via Vanta. Both moves reduce procurement friction for regulated buyers evaluating data annotation workflows. The remaining moves on Open Code Review, DeepSeek Reasonix, Qwen Code, OpenCursor, and OpenLIT were GitHub repository publications with limited additional detail in the pack.

WATCH

Two threads are worth tracking into the next cycle. First, [T4](#)'s open feature request on per-call cost settlement is a leading indicator: if merged, it would make Plano one of the few gateway-layer tools capable of supporting internal chargeback models for agentic workloads, a feature with direct relevance to any organization running multiple teams on shared AI infrastructure. Second, [T3](#) enters with strong compliance signals but no visible enterprise pricing tier yet; the freemium model and open-source deployment suggest a self-serve motion that may not yet be tuned for enterprise procurement. Watch for any commercial tier announcement. [T5](#) joining the AI2 Incubator also merits follow-up: incubator backing typically precedes a structured funding round and possible enterprise support tier.

Bottom line: This week's pack establishes a materially strengthened agentic stack for regulated buyers, with credible new entries across every layer from runtime enforcement to domain-specific agents, and 12 tracked-product moves that reduce procurement uncertainty. The single most important action is to run a structured comparison between [T5](#) and Glean before the next budget cycle: Apache 2.0 licensing, AI2 Incubator backing, 50+ connectors, and on-prem deployment combine to make PipesHub a cost-model disruptor in the enterprise search category, and the window to evaluate it before it attracts enterprise pricing is likely short.

PRODUCTS

PRODUCT	MATURITY	DEPLOYMENT	PRICING	FEATURES
LangAlpha Ginlix AI · A specialized version of Claude Code tailored for financial market ana ...	Open source project	saas, vpc, open_source	freemium	parallel research agents, sec filing analysis, live market data, options chains, treasury ...
LangChain LangChain · A general-purpose framework for building and orchestrating LLM applica ...	Open source project	saas, open_source	unknown	agent orchestration, agent observability, agent evaluation, agent deployment, agent fleet ...
Hoop Hoop · Open-source runtime enforcement tool that masks sensitive data and blo ...	Open source project	open_source, saas	freemium	pii masking, destructive command blocking, human approval routing, risk scoring of actions ...
Plano Katanemo · An AI-native proxy server and data plane for agentic apps providing sm ...	Open source project	on_prem, open_source	unknown	agent orchestration, model routing, guardrail hooks, jailbreak detection, agentic traces
PipesHub PipesHub · An open-source platform for securely connecting enterprise knowledge t ...	Open source project	open_source, vpc, on_prem, api	free	permission-aware search, explainable answers with citations, knowledge graph retrieval, 50 ...

Nuwax Agent OS Nuwax AI · An open-source enterprise-grade AI Agent Development and Operation Pla ...	Open source project	on_prem, open_source free	agent creation and distribution, knowledge base management, model proxy, memory system, pl ...
LlamaIndex LlamaIndex · A data indexing and retrieval framework primarily focused on RAG and c ...	Open source project	saas, api, open_source, freemium vpc	agentic document ocr, schema-based extraction, document indexing for rag, chart and table ...
Notte Nottelabs · A cloud browser infrastructure and web automation platform designed sp ...	Early stage startup	saas, api, open_source usage	cloud browser sessions, browser agents, serverless browser functions, credential vaults, a ...

MOVES ON TRACKED PRODUCTS

- PipesHub: news "PipesHub GitHub repository published with Apache 2.0 license, Docker-based deployment, and examples ..." on 2026-09-28
- PipesHub: news "PipesHub homepage self-identifies as 'The Open Source Glean Alternative' with 50+ connectors and MCP ..." on 2026-09-28
- PipesHub: news "AI2 Incubator LinkedIn post references PipesHub as open-source enterprise AI alternative joining the ..." on 2026-09-28
- LangAlpha: news "LangAlpha listed and described on aikii.org as open-source financial Agent OS" on 2026-09-28
- Open Code Review: news "alibaba/open-code-review" on 2026-09-28
- DeepSeek Reasonix: news "esengine/DeepSeek-Reasonix" on 2026-09-28
- Qwen Code: news "QwenLM/qwen-code" on 2026-09-28
- OpenCursor: news "PawanOsman/OpenCursor" on 2026-09-28
- OpenLIT: news "openlit/openlit" on 2026-09-28
- Plano: feature "Community feature request opens discussion on per-call agent cost settlement at the gateway layer" on 2026-09-25
- Encord: pricing "Encord pricing page updated to reflect VPC and on-prem add-on options alongside Starter, Team, Enter ..." on 2026-09-23
- Encord: news "Encord security page confirms SOC 2 examination completion and HIPAA compliance via Vanta" on 2026-09-23

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 2026-09-28.

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.