

The Agent Stack Meets the ROI Wall

This week the infrastructure pitch for enterprise AI got noticeably louder. AWS, Databricks, and Adobe each used their summer summits to stake out layers of what they describe as a full agent stack — governed data context, real-time analytics under the lakehouse, and creative agents embedded into the tools marketing teams already use. The announcements landed in the same week that surveys and venture-capital commentary made the counter-pressure plain: a growing share of generative AI pilots are being abandoned, and finance leaders are starting to demand evidence before signing the next renewal. The literacy question for this window is how to read vendor architecture moves against a buying climate that is rapidly losing patience.

TL;DR

- AWS used its NYC summit to announce Context, a self-learning knowledge graph for agents, alongside Continuum for agent security and an expanded Bedrock/Kiro/DevOps lineup — a bid to own the governed-data layer underneath enterprise agents.
- Databricks countered with Lakehouse//RT and LTAP, claiming up to 16x performance gains over dedicated real-time serving stacks and pitching a single data substrate for both human users and agents; performance figures are vendor- and customer-reported, not independently benchmarked.
- Adobe pushed its creative agent into public beta across Premiere, Photoshop, Illustrator, InDesign and Frame.io, and into ChatGPT, Claude, Copilot and Gemini — making AI-assisted production a default surface inside the dominant creative suite.
- Survey and analyst data this window point to a hard ROI turn: Gartner figures cited in coverage report more than half of GenAI proofs of concept abandoned and project 40% of agentic AI projects cancelled by end of next year, while a Genpact/HFS study of 2,000+ executives estimates ~\$18T in value blocked by data, process, tech and talent debt — a modeled figure, not an audited one.
- Export controls on a frontier model and the appearance of major AI CEOs at the G7 signal that AI sovereignty has moved into top-tier policy, with direct implications for who can access which models and where they can be deployed.

Three vendors, one pitch: own the agent stack

Within a 48-hour stretch, three large platform vendors each tried to define a different floor of the enterprise agent stack. AWS used its New York summit to introduce Context, a managed, self-learning knowledge graph designed to give agents governed access to enterprise data without manual curation, paired with Continuum for agent security and expansions to Bedrock AgentCore, Kiro and a DevOps agent. The framing across both AWS's own materials and outside coverage is consistent: agents are only as trustworthy as the data they can reach, and AWS wants that reachability to run through its services.

Databricks pushed at an adjacent layer. Lakehouse//RT, announced at its Data + AI Summit and covered alongside a new transactional/analytical processing approach, claims sub-100ms query latency at high concurrency and up to 16x performance gains over dedicated real-time serving systems, with named customers reporting multi-fold response-time improvements. The strategic argument is consolidation: collapse separate operational and analytical systems so agents and dashboards read from the same governed substrate. These are vendor- and customer-reported figures from a beta product, not independent benchmarks, and should be read as directional architecture signals rather than settled performance facts.

Adobe took the same logic into a very different surface. Its creative agent entered public beta across Premiere, Photoshop, Illustrator, InDesign and Frame.io, and was simultaneously wired into ChatGPT, Claude, Copilot, Gemini and Slack. Adobe's own survey of more than 16,000 creators, cited in its announcement, reports that 75% describe creative AI as integrated or essential to their work while 85% insist final creative decisions should remain human — a useful tell about how vendors are positioning agents as orchestrators of multi-step production rather than as replacements for judgment.

Read together, the three announcements describe a shared pitch: a context and governance layer (AWS), a unified data and serving layer (Databricks), and an embedded action layer inside the tools end users already open (Adobe). For technology and operations leaders, the practical question is less which vendor wins than how many of these layers a given organization wants under one roof — and what that implies for lock-in, integration cost, and the speed at which agents can be deployed against real workflows.

Sources: VentureBeat AI (<https://venturebeat.com/data/aws-enters-the-context-layer-race-with-a-graph-that-learns-from-agents-not-manual-curation/>); cio-dive (<https://ciodive.com/news/AWS-continuum-AI-security-claude-mythos/823184/>); aws.amazon.com (<https://aws.amazon.com/blogs/machine-learning/context-intelligence-for-your-data-and-ai-agents-at-scale/>); aboutamazon.com (<https://www.aboutamazon.com/news/aws/aws-summit-nyc-2026-ai-agents/>); databricks.com (<https://www.databricks.com/company/newsroom/press-releases/databricks-launches-lakehouse-rt-bring-real-time-analytics-directly/>); siliconangle.com (<https://siliconangle.com/2026/06/16/databricks-declares-end-pipelines-unified-platform-operational-analytical-data/>); news.adobe.com (<https://news.adobe.com/news/2026/06/adobe-unveils-major-expansion/>); blog.adobe.com (<https://blog.adobe.com/en/publish/2026/06/18/more-time-spend-on-your-craft-adobe-brings-power-creative-agent-creative-cloud-apps/>); 9to5mac.com (<https://9to5mac.com/2026/06/18/adobe-expands-firefly-capabilities-extends-agentic-tools-to-creative-cloud-apps/>)

The ROI wall arrives

While vendors expanded the stack, the demand side tightened. Coverage of analyst data this week reports that more than half of generative AI proofs of concept have been abandoned, and projects that as many as 40% of agentic AI initiatives will be cancelled by the end of next year. The same coverage notes that AI risk disclosures among large US public companies have risen sharply over two years, and points to debt-financed AI infrastructure and per-employee tooling costs as pressure points that boards and audit committees are starting to notice.

Venture-side commentary echoes the shift. A widely cited investor interview describes named enterprises burning through annual AI budgets in months and trimming model licenses as initial enthusiasm meets cost reality. A separate practitioner piece argues that most stalled programs fail not on the technology but on the basics: unclear ownership,

missing success criteria, and tools deployed before the business problem is defined. That diagnosis is vendor-adjacent — the author works for an automation platform — but the pattern it describes is consistent with the harder numbers elsewhere in the window.

A Genpact/HFS Research study of more than 2,000 executives, surfaced through both a vendor release and a bylined op-ed, puts a frame around the gap. It identifies four compounding 'enterprise debts' — data, process, technology and talent — and estimates roughly \$18 trillion in recoverable value across the Global 2000 if they are resolved, with potential for 8% faster revenue growth and 16% lower operating costs. The same study reports that only about a third of enterprise data is AI-ready, around a third of the workforce is AI-ready, and only 6% of enterprises have funded debt-resolution programs at scale. These are modeled figures from a vendor-sponsored study, not audited results, but the direction is hard to miss.

The contrast with the buying climate is sharper still. A separate CFO survey reports that two-thirds of finance leaders plan to increase technology and AI spending over the next year, even as economic confidence sits at a multi-quarter low, and that nearly all responding organizations are now piloting, scaling or fully integrating AI. The story of this window is not retreat — it is a bifurcation. Budgets are still moving, but the tolerance for open-ended experimentation is visibly narrowing, and the questions attached to each renewal are getting harder.

Sources: media.genpact.com (<https://media.genpact.com/2026-06-15-Enterprises-Are-Sitting-on-18-Trillion-in-Trapped-AI-Value-New-Research-Shows-How-to-Unlock-It>); [fastcompany.com](https://www.fastcompany.com) (<https://www.fastcompany.com/91559945/you-cant-build-your-ai-future-on-broken-foundations>); TechCrunch AI (<https://techcrunch.com/video/neas-tiffany-luck-says-enterprises-are-still-figuring-out-their-ai-roi>); bitcoinworld.co.in (<https://bitcoinworld.co.in/nea-tiffany-luck-enterprises-ai-roi/>); CFO Dive (<https://cfodive.com/news/cfos-boost-tech-spending-despite-economic-uncertainty-grant-thornton-ai/823210>); em360tech.com (<https://em360tech.com/podcasts/why-most-enterprise-ai-investments-fail-deliver-roi>); [techtarget.com](https://www.techtarget.com) (<https://www.techtarget.com/searchcio/feature/AI-market-correction-What-IT-leaders-must-know>)

AI gets a border

Two policy signals this week are worth holding alongside the vendor and ROI stories. The US administration imposed export controls on a frontier model and its underlying base model, restricting access for foreign nationals — including, per the reporting, those employed by the developer inside the United States. That is a notable extension of national-security tooling into the day-to-day mechanics of who can touch which weights, and it lands directly on hiring, deployment geography and customer-access questions that AI-using enterprises had largely treated as commercial decisions.

In parallel, leaders of the largest frontier-model companies joined heads of state at the G7, where the agenda included frontier AI risks, infrastructure, sovereignty and child safety. The presence itself is the signal: AI governance is now being shaped at the level where standards tend to harden into binding rules, with model developers at the table.

For executives, the practical literacy point is narrow but important. Vendor selection, workforce composition and cross-border deployment can no longer be treated as purely technical or commercial choices. The same agent stack being assembled by AWS, Databricks and Adobe will be deployed into an environment where which model is

allowed, by whom, and where, is increasingly a policy variable rather than a procurement one.

Sources: the-verge-ai-feed (<https://theverge.com/podcast/951542/anthropic-claude-fable-5-mythos-ban-pentagon-ai-regulation-trump>); cnbc.com (<https://cnbc.com/2026/06/17/g7-trump-ai-tech-leaders-openai-anthropic-google.html>)

Concept of the Week: The Agent Stack vs. the Evidence Stack

Vendors are assembling an 'agent stack' — context/knowledge graphs, real-time data engines, security layers, embedded copilots — and selling it as the foundation for production AI. Buyers increasingly need a parallel 'evidence stack': defined business problems, named owners, measurable outcomes, and exit criteria. When the two are out of sync, spend compounds faster than value. This window is the clearest sign yet that the two stacks are being built at very different speeds.

What to watch

Three threads to track into next week. First, whether any independent benchmarks or named customer outcomes appear for the AWS Context and Databricks Lakehouse//RT claims — beta-stage vendor figures are useful as direction, not as proof. Second, whether more enterprises publicly tie AI spend to specific outcomes, or quietly trim license counts the way several named companies reportedly have; the gap between rising CFO budgets and rising pilot-abandonment rates cannot stay open indefinitely. Third, how the export-control posture evolves: the restriction on foreign-national access to a frontier model inside the US is the kind of move that, if extended, would reshape hiring and deployment plans well beyond the companies directly affected.

Source Ledger

AWS enters the context layer race with a graph that learns from agents, not manual curation

<https://venturebeat.com/data/aws-enters-the-context-layer-race-with-a-graph-that-learns-from-agents-not-manual-curation>

AWS unveils agent security, data access tools

<https://ciodive.com/news/AWS-continuum-AI-security-claude-mythos/823184>

AWS launches 'Context Intelligence' stack to give AI agents governed access to enterprise data

<https://aws.amazon.com/blogs/machine-learning/context-intelligence-for-your-data-and-ai-agents-at-scale/>

AWS launches AI agent platform stack at Summit New York 2026

<https://www.aboutamazon.com/news/aws/aws-summit-nyc-2026-ai-agents>

Databricks Launches Real-Time Analytics Engine Directly on the Lakehouse, Eliminating Separate Serving Stacks

<https://www.databricks.com/company/newsroom/press-releases/databricks-launches-lakehouset-bring-real-time-analytics-directly>

Databricks bets on unified data architecture to eliminate the pipeline bottleneck blocking AI agents

<https://siliconangle.com/2026/06/16/databricks-declares-end-pipelines-unified-platform-operational-analytical-data/>

Adobe Rolls Out AI Creative Agent Across Photoshop, Premiere, Illustrator and Major AI Platforms

<https://news.adobe.com/news/2026/06/adobe-unveils-major-expansion>

Adobe embeds AI creative agent across Premiere, Photoshop, Illustrator, InDesign and Frame.io

<https://blog.adobe.com/en/publish/2026/06/18/more-time-spend-on-your-craft-adobe-brings-power-creative-agent-creative-cloud-apps>

Adobe Brings Agentic AI Across Creative Cloud, Firefly, and Major AI Platforms

<https://9to5mac.com/2026/06/18/adobe-expands-firefly-capabilities-extends-agentic-tools-to-creative-cloud-apps/>

\$18 Trillion in AI Value Trapped by Four Enterprise Debts, Study of 2,000+ Executives Finds

<https://media.genpact.com/2026-06-15-Enterprises-Are-Sitting-on-18-Trillion-in-Trapped-AI-Value-New-Research-Shows-How-to-Unlock-It>

85% of Enterprises Lack the Foundations to Deliver on AI Investments, Survey Finds

<https://www.fastcompany.com/91559945/you-cant-build-your-ai-future-on-broken-foundations>

NEA's Tiffany Luck says enterprises are still figuring out their AI ROI

<https://techcrunch.com/video/neas-tiffany-luck-says-enterprises-are-still-figuring-out-their-ai-roi>

Enterprise AI Spending Hits ROI Reckoning: Uber, Meta Pull Back as CFOs Demand Proof

<https://bitcoinworld.co.in/nea-tiffany-luck-enterprises-ai-roi/>

CFOs boost tech spending despite economic angst: Grant Thornton

<https://cfodive.com/news/cfos-boost-tech-spending-despite-economic-uncertainty-grant-thornton-ai/823210>

Why Most Enterprise AI Investments Fail to Deliver ROI

<https://em360tech.com/podcasts/why-most-enterprise-ai-investments-fail-deliver-roi>

AI Market Correction Warning: Half of Enterprise GenAI Projects Failing as Costs Outpace Returns

<https://www.techtarget.com/searchcio/feature/AI-market-correction-What-IT-leaders-must-know>

Trump Administration Imposes Export Controls on Anthropic's Fable 5 AI Model

<https://theverge.com/podcast/951542/anthropic-claude-fable-5-mythos-ban-pentagon-ai-regulation-trump>

'A signal of where power sits': Trump and world leaders joined by OpenAI, Anthropic, Google at G7

<https://cnbc.com/2026/06/17/g7-trump-ai-tech-leaders-openai-anthropic-google.html>

Corrections

No public corrections filed.

Production Metadata

anthropic/claude-opus-4.7 / generated Jun 18, 2026 / 18 sources cited.