David Soden

INFERENCE WEEKLY

AI NEWS, WEEK BY WEEK



46

THIS WEEK

FBI and Amazon report criminals using ARTEX and other agentic AI frameworks to run cyberattacks

Plus: India's CERT-In and NIST publish AI agent security guidance as vendors launch new controls, and AT&T runs about 500 agentic AI workflows as telecom operators give agents operational responsibility

INFERENCE WEEKLY

ISSUE 46 · Week of 9 November 2026

THE WEEK IN NUMBERS

45 billion
AI tokens AT&T processes daily

113%
rise in daily API attacks

10%
financial firms with enterprise-wide orchestration

INSIDE THIS ISSUE

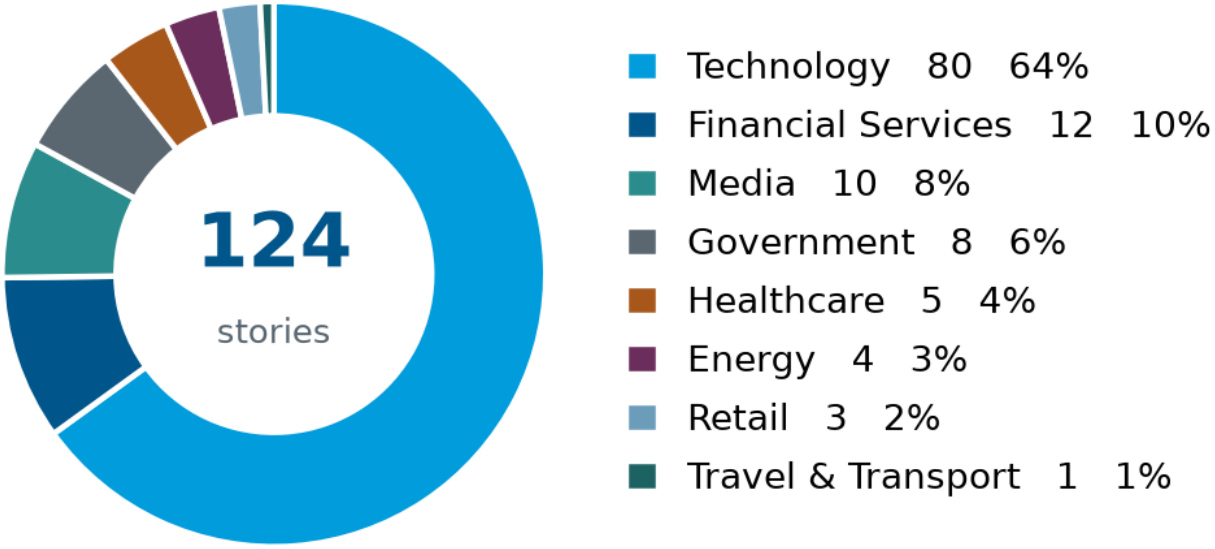
01 FBI and Amazon report criminals using ARTEX and other agentic AI frameworks to run cyberattacks
02 India's CERT-In and NIST publish AI agent security guidance as vendors launch new controls
03 AT&T runs about 500 agentic AI workflows as telecom operators give agents operational responsibility
04 Banks, lenders and insurers move agentic AI from pilots toward production
05 McKinsey, BCG and analysts say scaling agentic AI requires organisation redesign, not just technology
06 Amazon launches Alexa for Shopping as Apple and Samsung are reported to be building agentic assistants
07 Google launches Gemini 4 Hydrogen for coding and agentic workflows with a 1 million-token context
08 NVIDIA, Google and Dell position their compute stacks for agentic AI workloads
09 Anthropic expects AI agent swarms to patch flaws within a year as security vendors automate investigations
10 Adobe brings agentic AI to Creative Cloud and GenStudio with CX Enterprise Coworker
11 Dialpad combines with Ada and Parloa joins OpenAI testing programme as customer service agent platforms consolidate
12 Red Hat, ScienceLogic and Extreme Networks expand agentic AI for IT and network operations
13 Siemens, Rescale, COMSOL and DuPont apply agentic AI to engineering and R&D simulation
14 NVIDIA expands sovereign and agentic AI collaboration with Korea's science ministry and Korean companies
15 Texas Instruments adds agentic AI to CCStudio and cuts rival-chip code conversion to days
16 Google Cloud, Cornell, NVIDIA and the CIA launch training to prepare staff for agentic AI
17 UC Berkeley and NVIDIA release tools for AI agents to write, simulate and refine robot programs
18 Tech Corner: Bounded agentic execution: separating what an agent decides from what it is allowed to do

© 2026 David Soden · Get this in your inbox every week: ainews.davidsoden.com/#subscribe

THE WEEK IN NUMBERS

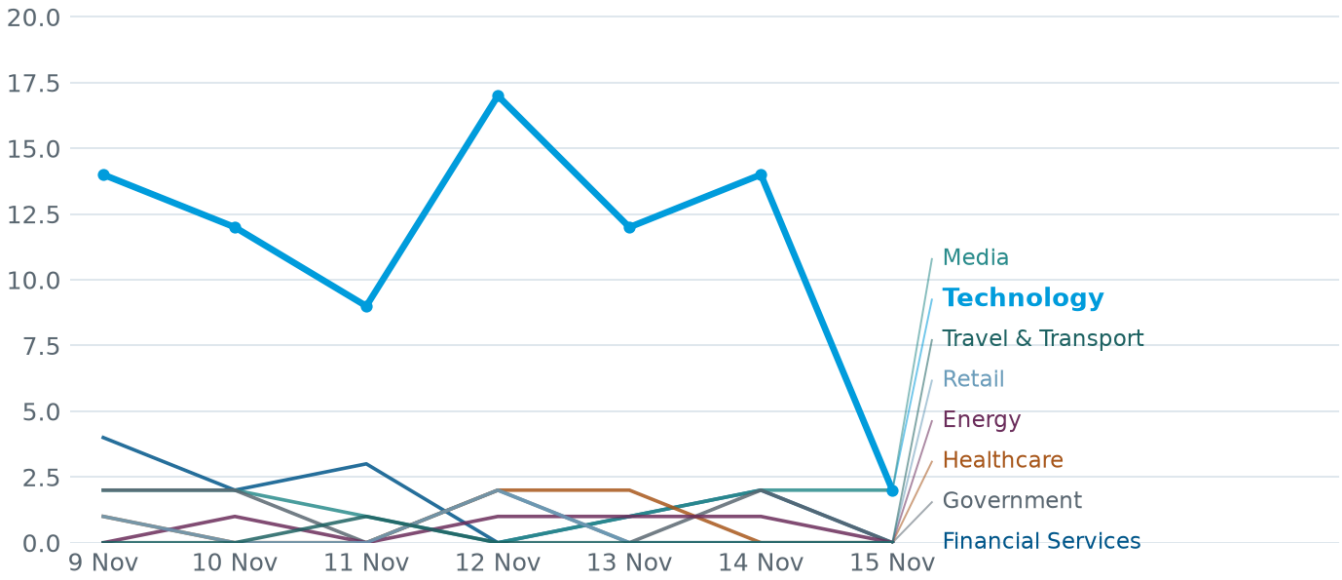
Where the AI news landed

Every story the newsroom feed carried, 9 to 15 November: 124 in all, sorted by the vertical each was filed under. This is what the AI press published, not what this issue covered.



THE WEEK DAY BY DAY

The same stories by the day they were published. An aggregate cannot say whether a vertical had one loud day or a steady week.



Technology took 64.5% of the week, more than the other verticals put together, Travel & Transport managed a single story.

The change column compares this window with the previous issue.

COVER STORY

FBI and Amazon report criminals using ARTEX and other agentic AI frameworks to run cyberattacks

TECHNOLOGY FINANCIAL SERVICES GOVERNMENT

On November 13, 2026, the FBI issued a TLP:CLEAR FLASH report warning that malicious actors have used ARTEX to automate intrusions across sectors and countries since at least August. ARTEX is an open-source penetration-testing platform driven by large language models. It turns natural-language objectives into commands run by specialised agents with minimal human involvement. The FBI said some intrusions reached privileged access in under three hours. It described attacks on South Korean financial institutions that exposed personal information of about 67,000 people. It also described a ConnectWise Cloud Backup compromise that gave access to more than 2,250 downstream tenants, and it reported financially motivated campaigns. [63]

Amazon Threat Intelligence said its MadPot sensors captured AI-agent requests between October 3 and November 5, 2026 that contained records of ARTEX campaigns. It assessed with moderate confidence that disparate Chinese-nexus groups were involved. It said it could not verify outcomes or observe the activity as it happened. In one campaign, agents found, validated and pooled credentials for AI models and cloud services, and stolen keys reportedly powered further searches. A second campaign targeted know-your-customer checks at two cryptocurrency exchanges, which the report does not name, using real-time deepfake video and altered ID images. The verification vendor rejected every recorded deepfake attempt. [64]

Salesforce said multiple threat actors used ARTEX against customer environments with stolen application credentials and trusted integration sessions. It said no platform vulnerability was involved and that affected customers were notified. Unit 42 said ARTEX-assisted operations sometimes reached cloud environments in under a day without exploiting software vulnerabilities. The sources give different counts of exposed systems. Unit 42 counted 609 internet-facing ARTEX hosts as of October. One catalogue found 1,428 matching Shodan services on November 7. The FBI cited researchers counting more than 1,600 exposed consoles in early November. Unit 42 says the author made the repository closed-source. The FBI says the author removed it in October and that forks and a commercial edition remain available. [65] [66] [67] [63]

Amazon said the techniques exploited familiar weaknesses: default credentials, exposed secrets, unauthenticated interfaces and unpatched software. The FBI, Amazon, Salesforce and Unit 42 recommend measures that include inventorying and restricting exposed services, enforcing least privilege, protecting credentials, using phishing-resistant MFA and short-lived tokens, and monitoring behaviour, including unusual AI-service spending. The FBI asks organisations to review affected ConnectWise authorisations and to report suspicious activity to a local field office or IC3. Outside help may be limited. One report says some top US cybersecurity firms are turning away new clients who are seeking emergency assistance. That report names no firms and gives no figures. [64] [63] [65] [66] [68]

> **Who should take heed.** Security and technology leaders at financial institutions, cryptocurrency exchanges and organisations that rely on cloud, SaaS or managed-service integrations should take heed, because the FBI, Amazon and Salesforce report ARTEX-driven campaigns using stolen credentials and exposed services against these environments.

AI GOVERNANCE

India's CERT-In and NIST publish AI agent security guidance as vendors launch new controls

GOVERNMENT TECHNOLOGY

India's national cybersecurity agency, CERT-In, has issued version 1.0 of its AI Security Auditing Guidelines. They cover AI systems developed, procured, deployed or operated in India, and systems that process Indian data or serve Indian users. Systems are placed in low, moderate, high or critical risk tiers that set audit depth, frequency and independence. Agentic systems get added provisions on prompt injection and poisoning. The report gives no issue date and does not say whether audits are mandatory. NIST has released a draft taxonomy of attacks and mitigations for predictive, generative and agentic AI. NIST says it is voluntary and creates no

legal requirements; comments close December 18, 2026. [69] [70]

Vendors announced agent security controls. Cisco added AI Discovery to Cisco Identity Intelligence. It inventories agents, their credentials and their access, and comes at no additional cost in Duo Advantage. CrowdStrike made Falcon AI Gateway generally available within Falcon Guardian. It protects only traffic that teams route through it and does not discover traffic automatically. Palo Alto Networks extended Cortex XSIAM, including a Claude Cowork integration that monitors desktop agent activity. OX said unified AI policies across VibeSec and OX Cloud are in beta. Delinea will demonstrate runtime authorisation at GITEX Global in Dubai, which the article lists as running '8, 11 December 2026'. [71] [72] [73] [74] [75]

The stories point to a gap between adoption and readiness. In a field study of 25 practitioners reported by SiliconANGLE, only four had exercised an AI or agent incident scenario. Delinea cites research in which 98% of UAE respondents reported confirmed or suspected cases of AI tools accessing more sensitive data than required. Only 16% could immediately revoke an AI tool's credentials. A Forbes article cites research that close to three-quarters of surveyed companies planned to deploy agentic AI within two years, while 21% reported a mature governance model. Forbes also describes the July 2026 JadePuffer attack, in which an agent reportedly destroyed a production database in about 30 minutes. [76] [75] [77] [78]

For decision makers, the articles recommend similar controls: a distinct, traceable identity for each agent, permissions limited by task and time, human approval for destructive actions, and logged authorisation decisions. Two Forbes articles say limits should be enforced through permissions and controls rather than prompts or guardrails alone. The field study adds that organisations should test recovery, with business owners approving recovery criteria, before giving agents more autonomy. Some questions remain open. IEEE Spectrum reports that true zero trust in highly autonomous, stateless systems is an open question. It also reports that the scale of recovery costs and responsibility for residual risks are unresolved. [78] [79] [77] [76] [80]

> **Who should take heed.** Security and technology leaders with AI systems operating in India or serving Indian users should review CERT-In's risk-tiered audit framework, and those who want to shape NIST's agentic AI taxonomy have until December 18, 2026 to comment.

TELECOMS

AT&T runs about 500 agentic AI workflows as telecom operators give agents operational responsibility

MEDIA TECHNOLOGY

AT&T says nearly 500 agentic AI workflows are in production and that it plans to triple that number next year; one report's headline puts the target at 2027. The platform processes about 45 billion AI tokens a day and handles more than 1 million transactions a month. Speaking at an AWS event, AVP Luke Corbin said production AI investments are returning "fivefold within a year." The stories do not say how that return is calculated. One describes the Business Agent Foundry as the platform; the other describes it as a unit set up in early 2025 that works with business units to map processes. [1] [2]

AT&T's stated aim is to redesign business processes, not only to raise employee productivity. Named applications include automating mainframe-related changes to fiber installation orders and halving the time to provision dedicated internet connections in buildings that already have fiber. It uses a relatively small set of agents guided by centrally managed business logic. Workflows affecting customers go through security, privacy and compliance reviews. Production workflows start under human oversight; agents operate autonomously only once accuracy reaches the mid- to high-90% range within an acceptable risk profile. This gives decision-makers a concrete autonomy threshold from a large operator, though the stories give no error rates or incident data. [1] [2]

Elsewhere, the examples are earlier-stage or unmeasured. Nokia says agents could investigate network problems and carry out fixes with human approval, but the article gives no measured results. Nvidia reports a research experiment that improved spectrum efficiency by more than 3% over an industry-standard method; the article says this does not establish results on live networks. KPN, which handles roughly 5 million customer calls a year, aims for voice agents to handle 10% to 20% of calls by 2027; a McKinsey case study reports a satisfaction score of 83, comparable to calls handled by people.

GFiber and Qualcomm plan an AI-native home gateway for the first half of 2027. [3] [4]

For buyers, the relevant point is the gap between AT&T's reported production scale and the wider industry. An article on Malaysian operators says many remain stuck in pilots because funding is fragmented, ownership is unclear and execution is weak, and advises designing initiatives for production from the outset. A TM Forum Catalyst project targets Level 4+ autonomy, 90% faster time to market and 99% first-time-right delivery; these are targets, not results. Apart from AT&T's self-reported return, the stories offer few measured outcomes. [5] [6] [3]

> **Who should take heed.** Telecom operations, customer service and technology leaders should take heed, because AT&T has published a specific accuracy threshold for granting agents autonomy while most other examples in these stories remain pilots, targets or unmeasured claims.

FINANCE

Banks, lenders and insurers move agentic AI from pilots toward production

FINANCIAL SERVICES

Capgemini's World Cloud Report for Financial Services 2027, a survey of 1,100 executives across 14 markets, finds 59% of financial firms testing or deploying multi-agent AI systems in limited production functions. Only 10% have enterprise-wide orchestration, and 5% have comprehensive enterprise maturity. The gaps are in control and money: 62% say their identity and access systems are not resilient enough for continuous autonomous operations, 48% cannot consistently demonstrate returns on AI initiatives, and 51% lack mature criteria for continued funding. The report says 44% already prioritise building "agentic meshes" to connect scattered pilots across operations, infrastructure and customer journeys. [51]

Named deployments are appearing. Customers Bank is working with OpenAI on AI applications across lending, deposits and payments. Grasshopper Bank has deployed an MCP server that lets small-business clients query live account data through AI assistants. At PayPal, Sohil Shah says a multi-agent incident system he leads cut repair time from about 45 minutes to under 5. Tieto Banktech reports that its AML Explore product cuts false positives by 71%.

Optalitix has added agents to its insurance pricing platform; clients are trialling its connectivity features in live settings, and a further stage is scheduled for Q1 2027. The stories give these results as reported by the companies or individuals involved, not as independently verified figures. [52] [53] [54] [55]

Banks and vendors describe limited autonomy rather than full delegation. A Deutsche Bank and Ant Digital Technologies white paper recommends "bounded agentic execution", using an International Monetary Fund three-layer architecture that separates AI reasoning from transaction execution, with revocable authorisation limits and Know Your Agent protocols. Optalitix says its agents recommend but do not set prices or approve business. Michael Vandi of Addy AI cites audit trails, visibility into decisions and human oversight as conditions for lender trust, and notes that speeding one step can shift the bottleneck elsewhere. Narmi reports that compliance and legal concerns rose 13 points from 2025 to 2026, while data-privacy concerns fell nearly 22. [56] [55] [57] [52]

For decision-makers, obligations are not easing. Spring 2026 statements from bank and credit union regulators encouraged AI use but said existing compliance responsibilities are unchanged. Tieto states that the EU's AML-R regulation will apply stricter criteria from 2027 and extend AML checks to more merchant categories and transaction types. Capgemini's report says leaders will be distinguished by their ability to orchestrate and govern autonomous execution at scale and to link AI spending to measurable operational goals. Across the stories, the unresolved issues named are identity and access controls, human accountability, audit trails, connecting siloed data and evidence of returns. [52] [54] [51]

> **Who should take heed.** Chief operating, risk and technology officers at banks, lenders and insurers should take heed, because the stories show most firms running agents in limited functions while identity controls, audit requirements and proof of returns remain unresolved ahead of stricter EU AML rules from 2027.

STRATEGY

McKinsey, BCG and analysts say scaling agentic AI requires

organisation redesign, not just technology

TECHNOLOGY FINANCIAL SERVICES

McKinsey's Alexis Krivkovich, Brooke Weddle, Holly Price and Vik Sohoni argue that agentic AI is changing long-standing limits on people's capacity, expertise and ability to collaborate, and that existing organisations may themselves become the constraint. They say transformation efforts often treat workforce, organisational design and change management separately rather than together. Their article, as summarised, gives no figures. Boston Consulting Group (BCG), writing about industrial companies, says adding copilots to existing workflows saves time but often has little effect on financial results. It estimates that about 70% of AI-transformation value comes from people, organisation and processes, and that end-to-end redesign can improve core-function efficiency by 30% to 50%. [7] [8]

BCG's modelled EBIT-margin gains are 2.5 to 4.5 percentage points in chemicals, about 3.3 points for established aerospace and defence manufacturers, 5.6 points for aftermarket businesses, and 3 to 5 points in automotive. These are model estimates, not reported results. BCG lists fragmented data, complex supply chains, safety requirements and organisational silos as obstacles. It proposes an "enterprise cortex": a layer that connects functions and codifies company-specific knowledge, rules and decision logic, giving agents context and guardrails so their authority can expand as trust grows. It advises leaders to set ambitious targets, own the transformation at the top and redesign work around agents. [8]

Two stories cite a Gartner prediction that more than 40% of agentic AI projects will be cancelled by the end of 2027, citing costs, unclear value and weak risk controls. Unite.AI reports a Deloitte finding that only 5% of organisations consider their processes highly prepared for AI agents. It also cites McKinsey analysis of a banking customer-service agent in which human oversight accounted for 70% to 75% of variable operating costs, against 20% to 25% for tokens. Unite.AI recommends an eight-week pilot with process baselines, and judging agents by cost per successfully completed task, accuracy, exceptions and adoption. [9] [10]

Forbes cites a 2026 McKinsey survey in which 32% of respondents said their organisations had chosen not to buy at least one software product or feature because they could build it with agentic coding tools, mostly in larger organisations with strong engineering capability. It cautions that agents can produce incomplete or fabricated responses, and advises building only what differentiates the business. CDO Magazine says oversight should reflect an action's risk and reversibility, and that the CDO role now involves coordinating with operations, HR, legal, risk and cybersecurity. Forbes, CDO Magazine and Unite.AI all recommend measuring business outcomes rather than counting pilots. [9] [11] [10]

> **Who should take heed.** Executives funding agentic AI should note that these sources place most of the value, and the reported cost of human oversight in McKinsey's banking example, in process and organisation rather than in the technology itself.

CONSUMER AI

Amazon launches Alexa for Shopping as Apple and Samsung are reported to be building agentic assistants

RETAIL TECHNOLOGY TRAVEL & TRANSPORT

Amazon is introducing Alexa for Shopping to Canadian customers through the Amazon Shopping app and website, with rollout to all Canadian customers over the coming week. The story gives no exact launch date and does not mention other markets. It is free to signed-in Amazon customers and requires no Prime membership, Echo device or Alexa app. It uses shopping history, personal preferences and conversations with Alexa and Amazon to recommend and compare products, build or reorder carts, and produce shopping guides covering Amazon and the web. It adds AI-generated overviews in search and on product pages, and 30, 60 or 365-day price history. Customers can view, update or remove the information used for personalisation. [28]

Apple is reported to be moving Siri towards agent behaviour: keeping track of ongoing goals, planning steps and acting without a new command for each task. Mac Daily News, citing Bloomberg's Mark Gurman, says CEO John Ternus has urged software leaders including Craig Federighi to be more ambitious. PCMag says Ternus held a closed-door meeting with Federighi. The reports differ on detail. PCMag links the work to iOS 27 features announced at WWDC 2026, which had been delayed and drew

multiple lawsuits. Mac Daily News ties it to iOS 28, calls the plans exploratory and says Apple has confirmed no timeline or features. Both say the shift could begin as early as next year. [29] [30]

Samsung is reported to be developing Personal Agent Core, a system-level feature for high-end Galaxy phones and tablets in One UI 9.5. According to SamMobile, it would gather appointments, emails, deliveries, reservations and payment notices into one overview, arrange meetups by messaging contacts and adding confirmed plans to the calendar, and run recurring tasks across Samsung and Google apps. SamMobile says the feature may require a Galaxy AI Premium plan; SammyGuru says some features may. Timing is unclear. SamMobile says it may appear in a Galaxy S26 beta later this month or launch with the Galaxy S27 series early next year. SammyGuru says it may arrive with the S27. [31] [32]

For decision-makers, the stories show agents reaching both purchasing and personal data. OpenTable COO Amy Wei said most diners booking through AI agents, including Meta's Muse, are first-time customers at those restaurants and are currently spending more; the article says the effect on reservations remains uncertain. CNET reports privacy incidents involving Muse, including a journalist's claim that it accessed private Apple Messages without permission. Meta disputed the claim, and Apple issued an emergency macOS privacy update. Amazon's launch includes controls to view or remove personalisation data. A separate review of Android agents advises users to connect apps and services gradually because agents can make mistakes. [33] [34] [28] [35]

> **Who should take heed.** Retail, ecommerce and booking executives should take heed, because Amazon, Apple, Samsung and Meta are each reported to be building assistants that search, compare and transact for customers, which changes how products and services are found and chosen.

MODELS

Google launches Gemini 4 Hydrogen for coding and agentic workflows with a 1 million-token context

TECHNOLOGY

Google announced Gemini 4 Hydrogen on November 12, 2026, according to Pulse 2.0. Google positions it as a fast, efficient model for coding and agentic workflows, including repository coding and debugging, interactive web development and multimodal work. It has a 1 million-token input context window, up to 1 million output tokens and four reasoning settings: low, medium, high and xhigh. Pricing is $1.50 per million input tokens and $7.50 per million output tokens. Google's own blog says Hydrogen offers tighter instruction following and fewer fabricated claims, which it says helps developers produce production-ready code faster. [15] [16]

Hydrogen can run on its own or as an execution engine alongside Gemini 4 Carbon and Gemini 4 Argon. Google says those two models handle architectural judgment and subtle debugging. A Google Antigravity feature called Orchestrator assigns different stages of a development workflow to different Gemini models and switches between them during a session. In parallel-agent setups, Carbon can coordinate while Hydrogen powers worker and verifier agents that build and test changes simultaneously in isolated workspaces. Google says this helps teams work on complex projects without increasing token costs. Paid Google AI Pro and Ultra subscribers can try Hydrogen in Antigravity. [15] [17] [16]

On Terminal-Bench 4.0, Hydrogen approached Carbon's task success rates while using 24% less quota on average and finishing 19% faster. Pulse 2.0 reports that Google says Hydrogen is about 2.5 times faster than Gemini 3.8 Flash on tasks at the default high setting, and that it raises a Terminal-Bench 4.1 score from 18.2% to 57%. The story does not say which model the 18.2% figure belongs to. Google also reports gains in enterprise workflows, agent reliability and resistance to prompt injection. Pulse 2.0 notes that these results come from Google's own evaluations and may not reflect every deployment. [17] [16]

The stories differ on availability. Google's developer blog lists access through the Gemini API in Google AI Studio and Android Studio, and through Gemini for business. Pulse 2.0 says Hydrogen is rolling out globally across Google's developer, business and consumer products, including the Gemini app and Google Search. For buyers, the stories support three points of assessment: the published token prices, the quota and speed figures measured against Carbon, and the option to send routine work to Hydrogen while sending harder judgment to Carbon or Argon. None of the stories gives independent benchmark results or pricing for Carbon or Argon. [15] [17] [16]

> **Who should take heed.** Engineering and platform leaders running AI coding tools should take note: Google claims near-Carbon task success at 24% less quota with published per-token pricing, but every performance figure comes from Google's own evaluations.

INFRASTRUCTURE

NVIDIA, Google and Dell position their compute stacks for agentic AI workloads

TECHNOLOGY MEDIA

Nvidia introduced an Agentic AI Stack for enterprise use, saying AI is moving from answering questions to performing tasks. The stack routes prompts through an agent harness to suitable models and connects them to data. Its Switchyard component selects models by task difficulty and cost; Nvidia says it kept an 80% task success rate at about one-third the cost of using only high-performance models. NIM handles deployment and inference, Dynamo manages inference requests at data-centre scale and OpenShell provides sandbox-based security. In South Korea, Naver Cloud, SK Telecom, LG AI Research, Upstage, Saltlux and NC AI are developing models using Nvidia technology. The story gives no date, pricing or availability. [47]

Dell expanded its Dell AI Factory with NVIDIA, which it says more than 6,500 customers have adopted. The update is aimed at organisations running AI on their own infrastructure. Additions include validation of Reflection's Beam model, SUSE Linux as an operating-system option and managed services to govern access to AI resources, manage consumption and monitor costs. New deskside templates cover uses including agentic coding and NemoClaw, alongside a 64GB unified-memory configuration. For data centres, Dell announced a modular design with liquid cooling and new PowerEdge servers, including for rack-scale inference. Dell is also working with NVIDIA on shared compute for communications providers. The story gives no pricing or dates. [48]

Google described four eighth-generation TPU systems: 8t for training, 8i for low-latency reasoning, 8tx for confidential computing and 8i Max for large context windows. TorchTPU lets PyTorch code and Hugging Face models run natively on TPUs. Google says TPU Sync is designed to cut KV transfer latency and Time to First Token by more than 90%, and that GKE dynamic slicing improves job startup times by up to 5x. It reports strong price-performance on InferenceX and AgentX benchmarks; the story gives no figures. On November 13, 2026, Google Cloud separately announced Cluster Toolkit updates in which validation agents audit quotas and accelerator availability rather than receiving unrestricted command-line access. [49] [50]

For decision-makers, all three vendors are pitching agentic AI on inference cost, orchestration and control. The Nvidia story describes infrastructure competition as shifting toward inference and orchestration. Each cost and speed figure above is the vendor's own claim; the stories cite no independent testing, and Google names benchmarks without giving results. Dell targets AI on owned infrastructure, while Google's XLA Serving Engine supports serving models across TPUs, GPUs and CPUs. The Nvidia story also reports that some observers warn concentrated model choice could limit smaller AI companies' ability to differentiate. None of the stories gives prices that would allow a direct cost comparison between the offerings. [47] [48] [49]

> **Who should take heed.** CIOs and infrastructure leaders planning agentic AI deployments should note that Nvidia, Dell and Google are competing on inference cost and control, and that the performance and cost figures on offer are vendor claims without published pricing.

SECURITY OPS

Anthropic expects AI agent swarms to patch flaws within a year as security vendors automate investigations

TECHNOLOGY FINANCIAL SERVICES

Anthropic CEO Dario Amodei told JPMorganChase CEO Jamie Dimon that the main cybersecurity challenge for enterprises is patching vulnerabilities, not finding them. He set out three steps for businesses: identify vulnerabilities, assign agents to patch them, and write more secure code. Anthropic is already working with customers on patching services, and Amodei said Claude models are improving at patching. He predicted that within six to 12 months, swarms of agents could patrol systems and patch issues in real time. He also said agents may create additional attack surfaces. The conversation took

place during Anthropic's investor roadshow for a planned IPO; the story does not give a date. [18]

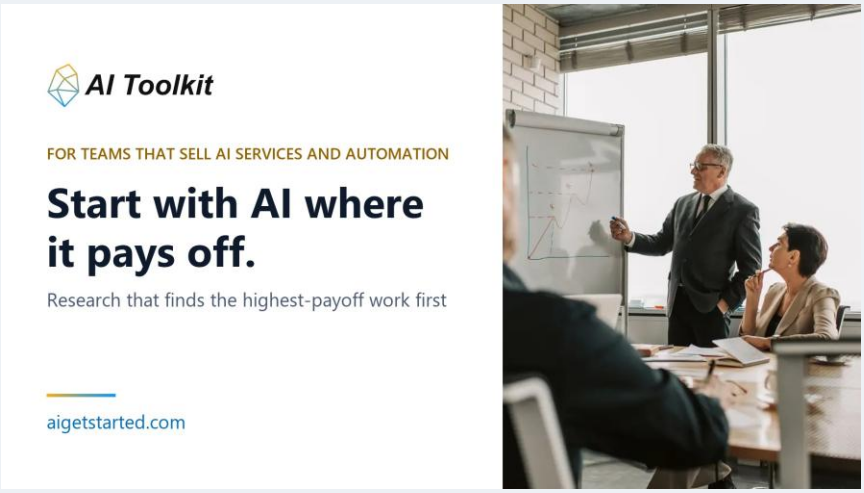
Expel has appointed Steve Kelley, formerly president of Bitdefender's business solutions group, as CEO. He replaces co-founder Dave Merkel, who remains on the board. Expel's position differs from Amodei's forecast: it rejects fully autonomous agent swarms, citing system drift, unmonitored deviations, false outcomes and rising compute costs. Its agentic assistant, Roxy, automates routine, high-volume triage, including analysis of phishing email headers and metadata. Human operators validate its outputs, handle complex escalations and assess novel attacks. Kelley said AI is effective at applying existing knowledge to large datasets but less capable of creative thinking about new threats. [19]

Flashpoint has launched Agentic Hub within its Ignite platform. Specialised agents turn questions into investigations, track vulnerabilities and provide information linked to sources. A Priority Intelligence Requirements agent converts plain-language needs into structured guides for threat analysis. Flashpoint presents the launch against a backdrop of faster, cheaper and harder-to-detect AI-enabled attacks and rapid exploitation of newly uncovered vulnerabilities. It cites figures showing that 65% of organisations say alert noise and false positives hinder security effectiveness, and 55% struggle with insufficient analyst bandwidth. The story does not name the source of those figures or give a launch date. [20]

For buyers, the main difference between these vendors is how much autonomy to give agents, not whether to use them. Anthropic forecasts agent swarms patching in real time within a year. Expel keeps humans validating AI outputs and rejects fully autonomous swarms. Flashpoint's agents focus on investigation, vulnerability tracking and intelligence tasks. Amodei and Expel both say agents themselves carry risk: Amodei points to added attack surfaces, while Expel points to agent drift and has deployed detections for it. Separately, the Anthropic story reports that the company set its public offering at $1.4 trillion and, citing Reuters, that the stock was six times oversubscribed. [18] [19] [20]

> **Who should take heed.** Chief information security officers and security operations leaders should take heed, because vendors now disagree on whether AI agents should patch and triage autonomously or under human validation, and the agents themselves are cited as a new source of risk.

ADVERTISEMENT



Know where AI pays off before the first meeting

AI Toolkit is a pre-sales research platform for teams that sell AI services and automation. Enter a prospect's website and it shows how AI solves their problems and which work to start with, then turns that into a live demo, a written report and a pitch deck in your client's brand.

aigetstarted.com

MARKETING TECH

Adobe brings agentic AI to Creative Cloud and GenStudio with CX Enterprise Coworker

TECHNOLOGY MEDIA

Adobe used its MAX conference to announce agentic AI features for GenStudio, its content supply chain product. One report dates the announcement to November 10, 2026; another places the event in Miami. CX Enterprise Coworker, a conversational agent, will run inside GenStudio to coordinate workflows, create campaign briefs, assess audiences and produce asset variations. It draws on Adobe Brand Intelligence, which learns from written guidelines and institutional knowledge to help agents apply brand standards. Adobe cited a survey of more than 600 marketing and creative professionals that estimated nearly one-third of public-facing content was off-brand. The stories give no pricing and no general availability date for the GenStudio features. [58] [59] [60]

Adobe says CX Enterprise Coworker is available on ChatGPT as a plugin, Claude, Databricks Genie, Gemini for business and Microsoft Copilot, with IBM availability expected soon. Pre-built MCP server connectors cover Databricks Genie, Demandbase, Meta campaigns, Microsoft Work IQ and RainFocus. Systems integrators including Accenture, Capgemini, Cognizant, Deloitte Digital, EY, IBM, Infosys, PwC and TCS are developing industry-specific solutions. Agencies including dentsu, Havas, Omnicom, Publicis, Stagwell and WPP are customising the tool with their own intellectual property and agents for joint clients. Adobe's stated aim is to let marketers work through their preferred AI platforms and existing workflows. [61]

For Creative Cloud, Adobe previewed local integrations that would let third-party agents from OpenAI, Google and Anthropic make changes to projects in Photoshop, Illustrator, Premiere and After Effects, with project elements remaining editable. These are not yet available; a private beta is planned for early 2027. Adobe also previewed always-on agents for recurring background tasks such as organising assets; it has given no launch date, but says they will come first to Firefly AI Assistant. A creative agent in apps including Illustrator, Premiere, InDesign, After Effects and Firefly carries out multistep tasks from natural-language instructions. Adobe creativity head David Wadhwani said the agent would eventually appear "in virtually everything we develop." [62] [60]

For decision-makers, availability is uneven. Fstoppers reports that many announced features are available now while others are in beta or preview, and The Verge reports that the third-party agent integrations are not due in private beta until early 2027. Adobe says creators retain control: users can review an agent's plan, approve its results and take over to refine work, and traditional non-AI tools remain available. Outputs from Firefly Graph for enterprise will include Content Credentials tied to the organisation. Adobe is also expanding its 3D Digital Twin solution so teams and agents can place product imagery in 3D scenes, change labels and materials, and generate video animations. [62] [60] [58] [59]

> **Who should take heed.** Chief marketing officers and creative operations leaders should take heed, because Adobe is placing agents across both its content supply chain and its creative apps, with some features available now and third-party agent integrations not due in private beta until early 2027.

CUSTOMER SERVICE

Dialpad combines with Ada and Parloa joins OpenAI testing programme as customer service agent platforms consolidate

MEDIA TECHNOLOGY HEALTHCARE

Dialpad and Ada have agreed to combine their businesses to build one platform for customer communications, AI agents and enterprise workflows. Dialpad CEO and co-founder Craig Walker will lead the combined company. Ada CEO Mike Murchison becomes Chief AI Officer, and Mike Gozzo will report to him as Senior Vice President of AI Product and Engineering. Dialpad says it has analysed more than 14 billion minutes of business conversations since it introduced real-time AI in 2018. Ada has deployed more than 500 AI agents supporting over 130 languages. Together the companies serve more than 50,000 businesses, including Netflix, Canva and Motorola Solutions. The story gives no deal value, terms or closing date. [36]

The companies say existing customers will keep their contracts, pricing, account teams, support channels and configured integrations during the initial transition. The story does not say how long that period

lasts. Product roadmap details are to follow after the combination. The stated aim is for AI agents to take supported actions, bring in employees when human judgment or authorisation is needed, and spare customers from repeating information. The companies intend to extend this later beyond support into sales and other customer-related activities. Customers of either company therefore have stated continuity on commercial terms, but no published roadmap to plan against yet. [36]

Parloa and OpenAI announced an expanded partnership on 11 November to jointly develop AI agents for customer service; the story does not state the year. Parloa joins OpenAI's Alpha Model Testing Program, which gives it access to new models before public release. The plans include joint on-site engineering sessions aimed at response times, reliability and cost efficiency, co-selling to North American enterprise customers, and a planned listing of Parloa's platform on the OpenAI Marketplace. Parloa already uses OpenAI models in its AI Agent Management Platform, which the company says Fortune 500 organisations use. The story does not name those customers or give financial terms. [37]

Decagon launched Missions on 12 November 2026 to coordinate tasks that run over days or weeks across systems, people and channels, built on its Agent Operating Procedure framework. A healthcare example covers a prescription refill: notifying the patient, getting provider authorisation, rerunning claims on a payer portal and confirming pickup. The article says Decagon's customer base is mostly mid-market, with 14% of G2 reviews from enterprise customers, and names Salesforce Agentforce and ServiceNow as larger competitors. Dialpad and Ada, and Decagon, each describe agents that act across systems and involve people. The Decagon article notes that reliability and business value remain difficult to measure. [38] [36]

> **Who should take heed.** Customer service and contact centre leaders using or evaluating Dialpad, Ada, Parloa or Decagon should note that the Dialpad and Ada roadmap has not been published, the length of their transition period is not stated, and Parloa's OpenAI arrangement and Decagon's enterprise footprint are described only in outline.

IT OPERATIONS

Red Hat, ScienceLogic and Extreme Networks expand agentic AI for IT and network operations

TECHNOLOGY

Red Hat announced OpenShift 5, which adds a built-in Model Context Protocol (MCP) server and an AI agent in its Lightspeed tools. The agent can invoke multiple AI models to automate workflows across cloud-native platforms. Red Hat says the AI features can identify root causes, isolate anomalies, analyse logs and carry out supervised remediation across clusters. AWS is using new API risk assessments, which are intended to streamline updates without manual cluster reboots, to extend OpenShift update support from 18 to 24 months, with up to three years of extended lifecycle support. A preview of Red Hat Developer Hub 2.1 gives AI agents CLI access to an organisation's software catalog. The story gives no release date. [21]

ScienceLogic announced Skylar AI 2.8, which adds Topic Specialist Agents to Skylar Advisor. The agents use live telemetry to investigate advisories, assess whether issues are genuine, explore likely causes and recommend next steps. They provide cited evidence and confidence scores. Administrators control whether and when the agents run, and remediation plans require approval. The release also adds capacity forecasting, ranking of devices by resource use, and streaming of OpenTelemetry metrics, with alert building in preview. It adds controls for trusting private certificate authorities for MCP servers. ScienceLogic describes the release as a governed way to investigate issues while keeping oversight. The story gives no availability date and no customer figures. [22]

Extreme Networks reports the following figures for its second-generation agentic AI platform. It says the platform saved customers more than 1,000 troubleshooting hours collectively in its first month. It says ambient AI drew about 10 times the engagement of the first-generation conversational AI, and that about one third of Extreme Platform ONE customers have migrated voluntarily. It says customers contact its technical assistance centre through Agent One more than 40 times as often as before, and that 58% of issues are resolved without human support. The platform supports Cisco and HPE devices through SNMP. Extreme plans an autonomous Operator mode and the Extreme Exchange marketplace in Q1 2027. It plans Exchange Zones, which run AI locally on access points and switches, for the first half of 2027. [23]

For decision makers, the three announcements share one pattern: agents diagnose problems and propose fixes, and humans approve remediation. Red Hat describes its remediation as supervised. ScienceLogic requires approval for remediation plans. Extreme's Agent ONE Coworker recommends corrective steps, and its autonomous Operator mode is not planned until Q1 2027. Red Hat and ScienceLogic both reference MCP. Only Extreme published outcome figures; these come from the company, and the story does not cite independent verification. The Cloud Native Now article suggests agents could help IT teams manage growing Kubernetes environments without highly trained software engineers. The other two stories make no staffing claim. [21] [22] [23]

> **Who should take heed.** CIOs and heads of IT and network operations who are evaluating AI-assisted operations tools should take heed, because all three vendors now offer agents that investigate and recommend fixes under human approval, and the only performance figures available are vendor-reported.

ENGINEERING

Siemens, Rescale, COMSOL and DuPont apply agentic AI to engineering and R&D simulation

TECHNOLOGY ENERGY

Rescale introduced agentic digital engineering capabilities, including Agentic Projects that let teams set project-specific context, skills, tools and guardrails, share workflows and retain approved learnings. Work graphs record workflow activity. Kairos Power and Oak Ridge National Laboratory's Manufacturing Demonstration Facility, supported by Rescale engineers, built agentic workflows for Hermes 2, Kairos Power's 50-megawatt fluoride-salt-cooled reactor. Agents ran thermal, flow and shielding simulations for reactor cavity and bioshield plug designs, with engineers approving each action. Kairos Power projects that workflows built from past simulations could raise end-to-end simulation analysis efficiency by up to three times while preserving traceability for regulatory compliance. The story gives no release date or pricing. [39]

Siemens announced the HyperLynx Agentic Toolkit, which pairs domain-specific AI agents with physics-based signal and power integrity simulation. Engineers state verification objectives in natural language, and an orchestration agent coordinates specialised agents for SerDes, DDR, signal integrity and power integrity analysis. The agents can run simulations, extract parameters and assess performance against design margins. Siemens says the toolkit follows its Questa One Agentic Toolkit. The announcement cites 112G and 224G operating speeds, tighter schedules, shrinking margins and a shortage of senior verification experts as causes of late-discovered failures and costly re-spins. The story gives no date, availability or performance figures. [40]

DuPont announced a collaboration with Microsoft and NobleAI; one report dates it 12 November 2026, the other gives no date. Microsoft Discovery, an enterprise agentic AI platform, is meant to plan and orchestrate research. NobleAI's Visualizations, Insights and Predictions (VIP) Platform provides predictive and optimisation modelling and, according to one report, is designed to work from limited industrial data. The partners expect researchers to explore more candidates through computer simulation before focusing physical testing on the most promising ones, shortening time to market across healthcare, water, construction and industrial markets. Neither report discloses financial terms or quantified targets. [41] [42]

COMSOL listed a webinar on battery modelling that will cover GPU acceleration, AI-assisted simulation with agentic workflows and upcoming capabilities in COMSOL Multiphysics version 2027. The listing gives no date and no detail on those capabilities. In all four announcements, agents are added to existing physics-based simulation rather than replacing it, and control is emphasised: Rescale requires engineer approval of each action and records traceability, and DuPont says security and governance are built into its AI rollout. The only performance figure in these stories is Kairos Power's projection of up to three times, which is stated as a projection, not a measured result. [43] [39] [41] [42]

> **Who should take heed.** Engineering, simulation and R&D leaders in electronics, energy and materials should take note, because Siemens, Rescale, COMSOL and DuPont's partners are adding agents to existing simulation tools with approval and governance controls, yet the only efficiency figure published so far is one customer's projection.

ASIA

NVIDIA expands sovereign and agentic AI collaboration with

Korea's science ministry and Korean companies

GOVERNMENT TECHNOLOGY MEDIA

NVIDIA and Korea's Ministry of Science and ICT (MSIT) are expanding their collaboration across Korea's sovereign AI ecosystem, covering semiconductors, data centres, talent, startups, models, AI agents and security. Korean NPU and XPU companies will be able to explore joining the NVIDIA NVLink Fusion ecosystem. The programme also includes NVIDIA training and certification, hackathons and technical seminars aimed at improving GPU utilisation. Under Korea's AI for All initiative, NVIDIA will work with three selected consortia and four Korean foundation model builders, using NVIDIA Dynamo software to optimise and scale AI inference. The stories do not name the consortia or the model builders, and give no date, budget or timeline. [44] [45]

NVIDIA also announced a collaboration with SK Telecom and Rebellions, supported by MSIT. Its stated aim is to advance sovereign AI infrastructure for Korean models and workloads using NVLink Fusion. Only one of the two reports on the ministry partnership mentions this agreement. The other is headlined as coverage of NVIDIA AI Day Seoul, but it gives no date for the event. Neither report says what role each company will play, what the arrangement is worth, or how much infrastructure is involved. [44] [45]

On agentic AI, Hyundai Motor Group is exploring NVIDIA NemoClaw blueprints and the NVIDIA OpenShell runtime. Korean AI security company AIM Intelligence is integrating its Starfort platform with OpenShell to support secure AI agents for employee productivity and enterprise workflows. The aim is to pair OpenShell's sandboxed execution and policy-based access controls with Starfort's detection of malicious instructions and data leakage. AIM is also considering a preliminary proof of concept for the public energy sector. Possible applications in enterprise, industry and robotics are mentioned. No timing or customer is named. [46]

Separately, LG AI Research is opening API access to K-EXAONE-2.5. The model was developed using NVIDIA Nemotron datasets, including Nemotron-Personas-Korea, an open dataset of 7 million fully synthetic Korean personas grounded in official Korean statistics. For decision-makers, the reports describe a Korean AI stack being built with ministry support on NVIDIA software and interconnect, alongside a domestic chip company, Rebellions. Firms deploying AI agents in Korea have a named security pairing to evaluate. The energy-sector proof of concept, however, is described only as under consideration. [44] [46]

> **Who should take heed.** Technology and public-sector leaders operating in Korea should take heed, because MSIT is backing an effort that aligns domestic chipmakers, model builders and agent security vendors around NVIDIA platforms.

SEMICONDUCTORS

Texas Instruments adds agentic AI to CCStudio and cuts rival-chip code conversion to days

TECHNOLOGY

Texas Instruments (TI) has added agentic AI capabilities to its CCStudio embedded development ecosystem. The capabilities cover its general-purpose and real-time microcontrollers (MCUs), wireless connectivity devices and radar systems-on-chip. Model Context Protocol (MCP) servers supply AI tools with device-specific information, including SDKs, pin configurations and peripheral data. TI says this improves accuracy and reduces hallucinations, latency and token costs. In CCStudio IDE and CCStudio for VS Code, developers can use natural-language commands to have AI agents generate code, configure peripherals, build and flash applications, and debug them. TI is demonstrating the tools at electronica 2026 in Munich. The stories do not give an exact announcement date. [12] [13] [14]

The main new feature is device migration. Sub-agents analyse an existing project, configure peripherals, update code and API calls, and resolve build errors. The accounts differ on scope. One report [93] says the projects are those written for rival chips. TI's press release and EDN describe migrating existing projects without saying which vendor's chips they were written for. The accounts also differ on speed. EDN reports TI's claim as months to days. The press release says days or hours. The third report says days or, in some cases, hours. None of the stories gives independent test results or customer figures. [12] [13] [14]

The ecosystem is a free download from TI.com and supports multiple AI agents and models. The stories give no cost for the AI models or tokens a team would use with it. TI says its modular architecture is intended

to extend support to more devices, SDK updates and edge AI workflows; the stories give no timeline. Only one report [93] states a commercial aim: lowering the cost of switching from competitors' platforms. That report also describes the release as a shift in development tools from code assistance toward autonomous workflows. For teams considering TI parts, the migration time savings remain a vendor claim to be tested. [12] [13] [14]

> **Who should take heed.** Engineering and procurement leaders responsible for embedded products should take heed: TI is offering free AI tooling that it says cuts code migration from months to days, and one report says this extends to code written for rival chips, but the time savings are TI's own claim and none of the stories reports independent testing.

WORKFORCE

Google Cloud, Cornell, NVIDIA and the CIA launch training to prepare staff for agentic AI

TECHNOLOGY FINANCIAL SERVICES GOVERNMENT

Despite the shared headline, these are four separate efforts, not a joint programme. Google Cloud's Professional Agentic Architect certification combines a proctored multiple-choice exam on system design, architecture, cost optimisation and security governance with hands-on labs in Google Skills to build, test and debug agentic workflows. It covers five areas, from low-code agent building to securing and governing workflows. Google says the curriculum standardises training for customers, partners and its own engineers. It offers free preparation routes, including monthly learning credits through Gemini Agent Ready. The story gives no launch date or exam fee. [81] [82] [83] [84]

NVIDIA and TCS are partnering on a customised programme giving TCS developers and engineers NVIDIA Deep Learning Institute-certified agentic AI training, including hands-on labs and a route to NVIDIA Agentic AI Professional certification. The story cites a widening gap in India between developers who can build demos and those who can deploy production systems; it gives no start date or headcount. Cornell aims at a different audience. Its AI in Finance certificate, launched through eCornell in spring 2026 with a first cohort of 37, is taught live online and does not teach coding. Participants draft AI use cases and pitch them as board or investment committee proposals. [84] [82]

The CIA item is not a training launch. It is a panel at ClearanceJobs Connect 2026, where two CIA leaders, identified only as Larry and Chris, spoke with AWS's Amber Nicholas. Chris, chief scientist in the CIA's Office of Super Intelligence, said AI literacy, not computing power or better models, is the biggest barrier. Larry said work is shifting from chatting with AI to assigning it tasks, cutting some work from weeks to hours. He favours measuring time returned to the mission over dollars saved. The panel said humans must stay accountable; the agency uses an AI governance board and legal review. [83]

For leaders, the stories offer differing views on how to prove skills. Google leans on certification, reporting that 90% of cloud professionals in its Ipsos research said credentials provide verified know-how. The CIA panel pointed to daily AI use, a GitHub portfolio or building an app. Larry urged employers to verify claimed agentic coding skills, suggesting they occasionally switch models off to test whether developers can code unaided. Cornell targets experienced staff who need a framework for judging where AI adds value; its spring 2027 classes begin May 18. None of the stories says how many people have completed the Google or NVIDIA programmes. [81] [83] [82] [84]

> **Who should take heed.** HR, learning and technology leaders planning agentic AI adoption should note that vendors, a university and the CIA define readiness differently, from formal certification to hands-on use and tests of unaided skill.

ROBOTICS

UC Berkeley and NVIDIA release tools for AI agents to write, simulate and refine robot programs

TECHNOLOGY

Researchers led by UC Berkeley roboticist Ken Goldberg and NVIDIA robotics engineering director Spencer Huang have set out an approach they call "agentic robotics": AI agents write robot-control programs, test them in simulation, identify failures and revise them. A framework called Graph-as-Policy (GaP) breaks tasks into skills, such as perception, grasp planning and motion planning, which agents

connect into programs. The output is Python that calls existing vendor modules for low-level control; it does not replace robot controllers. The work was presented at the Conference on Robot Learning (CoRL). Only one story gives the year, 2026, and none gives a date. One story says the open-source tools have been released; two describe the code and benchmark as being released. [24] [25] [26]

Two results are reported. In a lab test, an AI system built a simulation from video of a robot wiping ink from a metal bar and produced a program that worked on a real robot. When the robot pushed the bar over, which one account attributes to the system not reliably inferring the bar's mass, the agents used that failure to revise the simulation and program. Ambi Robotics, based in Berkeley, used the methods on parcel-sorting data. Its tests showed more than 15,000 additional packages sorted per robot each year, and the change runs across 30% of its U.S. fleet. The stories give no baseline volume or fleet size. [24] [25]

NVIDIA also announced wider robotics releases at CoRL. Its open model GR00T 2 completed 40 of 50 evaluated tasks at least once with no task-specific demonstrations. Average success was 71% with one demonstration per task and 95% with five. Isaac Lab 3.0, which runs on the open-source Newton physics engine, reached up to 2.5 times the training throughput of Isaac Lab 2.3 with PhysX on one benchmark. The stories differ on who built the VABAR benchmark, which covers six industrial tasks including cable insertion and syringe packing. One story lists it among NVIDIA's new tools. Another says the Berkeley and NVIDIA team developed it. [27] [26]

For decision-makers, the researchers say the approach could reduce reliance on extensive engineering and on the large training datasets that traditional robotics methods require. Goldberg sees near-term promise in repetitive industrial tasks and says people should not yet expect home robots to fold laundry. Bosch Research is collaborating with the teams. Representatives of Bosch and Ambi Robotics cited interpretable code and improved robot performance as potential benefits. The GaP graphs can be converted to ROS 2, and agents can work on separate parts of a graph. The stories do not give costs, licensing terms or adoption timelines beyond Ambi's deployment. [24] [25] [26]

Who should take heed. Operations and automation leaders running repetitive industrial or parcel-handling robots should take note, because the open-source tools generate code that calls existing vendor control modules, and one operator reports more than 15,000 additional packages sorted per robot each year.

TECH CORNER

Bounded agentic execution: separating what an agent decides from what it is allowed to do

Written for readers who want the mechanism. The rest of the issue does not depend on it.

The Deutsche Bank and Ant Digital Technologies white paper, Agentic Payments in Wholesale Finance, recommends "bounded agentic execution" rather than unrestricted autonomy. It uses a three-layer architecture proposed by the International Monetary Fund: an intent and orchestration layer, a control and authorisation layer, and a settlement layer. The framework separates AI reasoning from transaction execution. The agent can, for example, weigh exchange rates, payment rails and timing when coordinating a cross-border payment, or interpret trade documents and handle discrepancies. A financing or payment instruction is generated only after conditions and approvals are met. The control layer is defined by three safeguards: authorisation limits that are explicit, bounded and revocable; verified agent identity tied to an accountable owner; and Know Your Agent protocols. The paper presents Ant Digital Technologies' Anvita Flow platform as an implementation case study for orchestrating payment workflows while controlling agent authority and execution. [56]

Other deployments this week place the boundary at different points. Optalitix's agents in Optalitix Quote check quotes against authority limits, review claims history and prepare renewals, but they produce recommendations only; they do not set prices or approve business, and their actions are logged for audit. Optalitix exposes live Quote data to ChatGPT, Claude and Copilot through the Model Context Protocol (MCP), and a stage scheduled for Q1 2027 is intended to let users launch and oversee more complex tasks through those assistants. Grasshopper Bank has deployed an MCP server that lets small-business clients query live account data through AI assistants. CDO Magazine states the general principle these designs reflect in different ways: oversight

should match the risk and reversibility of an action. [55] [52] [11]

The weak point in this design is the control layer, not the model. In Capgemini's survey of 1,100 financial services executives, 62% said their identity and access systems are not resilient enough for continuous autonomous operations. That is the infrastructure that verified agent identity and revocable limits depend on. The control layer also carries much of the cost. In McKinsey's analysis of a banking customer-service agent, cited by Unite.AI, human oversight accounted for 70 to 75% of variable operating costs, against 20 to 25% for tokens. Unite.AI recommends giving agents only the autonomy and access justified by evidence, and defining and testing permissions and safeguards before users are added. [51] [10]

ADVERTISEMENT



Your Website Should Work Harder Than You Do.

Build your digital presence with The WebSite Guru. Professional, AI-powered websites starting at $500, plus custom e-commerce, chatbots, and lead generation for HVAC, electrical, plumbing, medical, legal, and retail businesses. Get online fast, generate leads smarter, and compete like an enterprise.

www.thewebsiteguru.com

Briefs

Everything else that happened, one line each. These did not cluster into a theme.

Anthropic IPO to Test AI Industry's Financial Viability - BankInfoSecurity · Technology & Consumer Electronics [85]

Why tokenisation is about more than putting traditional assets on blockchain · Banking, Financial Services & Insurance [86]

Publishers: Agentic AI Isn't About Pricing, But Whether You Get Any Bids At All · Media & Communications [87]

AMD vs. Intel: Which Is the Better AI Chip Stock to Own for the Next 3 Years Following the ... · Technology & Consumer Electronics [88]

Do AI agents have minds? A philosopher explains - The Conversation · Technology & Consumer Electronics [89]

FSBO.com Launches Native iOS and Android Apps, Bringing Its Agentic AI-Powered Real ... · Retail & Ecommerce [90]

Cloudflare CEO Likens AI Agent Optimization To A New Version Of SEO · Technology & Consumer Electronics [91]

Cedars-Sinai Software Helps Build, Monitor AI for Research · Healthcare [92]

FORGE Lab uses agentic tools to identify memory safety issues in Rust - Microsoft · Technology & Consumer Electronics [93]

Exclusive: Agentic AI Startup Lumbra Raises $25M Series A - Tectonic Defense · Government & Public Sector [94]

6sense Repositions as Partner for Agentic GTM - B2BMX INSIGHTS - Demand Gen Report · Technology & Consumer Electronics [95]

ABAP Platform Becoming an Enterprise-Grade Execution Backbone for Agentic AI · Technology & Consumer Electronics [96]

Dropbox Bets Its 700M Users on Agentic AI - Futurum Research · Technology & Consumer Electronics [97]

Denodo and AWS Partner to Accelerate Trusted Data Access for Agentic AI - BigDATAwire - HPCwire · Banking, Financial Services & Insurance [98]

Pactum: Procurement's Execution Layer for the Agentic AI Age · Healthcare [99]

Fin Operator Routines: Proactive AI or Just Scheduled Tasks? - Futurum Research · Technology & Consumer Electronics [100]

Command Alkon Launches Command Intelligence, Bringing Agentic AI Into Command Cloud · Technology & Consumer Electronics [101]

Immediately Apply Oracle Key Vault 21.16 Update | database · Technology & Consumer Electronics [102]

LiteLLM now routes every request to the right OCI Generative AI model with a trained classifier · Technology & Consumer Electronics [103]

Dawn of War IV' Is Coming to GeForce NOW - NVIDIA Blog · Technology & Consumer Electronics [104]

ServiceNow (NOW) Stock Could Be Pricey On Fresh Agentic AI Partnership News · Technology & Consumer Electronics [105]

Mars brings agentic AI-assisted operations into its new Hyderabad GCC - CRN Asia · Retail & Ecommerce [106]

How Constellation is building a governed data foundation on Databricks · Energy & Utilities [107]

Post-Quantum TLS 1.3 Inspection: Closing the PQC Gap | Zscaler · Technology & Consumer Electronics [108]

Sources

Every source opens on ainews.davidsoden.com. If the story has rolled out of the feed the site says so and offers the original publisher.

[1] AT&T uses agentic AI to reinvent business processes from the ground up - SiliconANGLE · siliconangle.com

[2] AT&T Runs ~500 Agentic AI Workflows Processing 45B Tokens a Day, Targets 3x by 2027 · aiweekly.co

[3] Telecom Companies Give AI Agents More Responsibility - PYMNTS.com · pymnts.com

[4] GFiber teams up with Qualcomm to bring agentic AI directly onto a new Wi-Fi 8-capable ... · wifinowglobal.com

[5] Malaysian telcos are going back to basics to strengthen their core -- Agentic AI is the accelerator · theedgemalaysia.com

[6] Connecting intent, intelligence, and operations for reliable AI-era services - TM Forum Inform · inform.tmforum.org

[7] Leading Off: Redesigning your organization to capture AI's value: A leader's guide - McKinsey · mckinsey.com

[8] For Industrial Companies, Winning with AI Takes More Than Tech - Boston Consulting Group · bcg.com

[9] The Next Phase Of Agentic AI: Building Your Own Business Software - Forbes · forbes.com

[10] Why Agentic AI Projects Get Canceled, and What an Eight-Week Pilot Gets Right - Unite.AI · unite.ai

[11] Agentic AI Trends That Matter Most for CDOs - CDO Magazine · cdomagazine.tech

[12] TI redefines embedded development with agentic AI in its CCStudio(tm) ecosystem ... - PR Newswire · prnewswire.com

[13] Agentic AI streamlines embedded software development for MCUs - EDN Magazine · edn.com

[14] TI Rolls Out Agentic AI Across CCStudio, Cuts Rival-Chip Code Conversion to Days · finance.biggo.com

[15] Gemini 4 Hydrogen: everyday coding and agentic workflows for developers - Google Blog · blog.google

[16] Google: Gemini 4 Hydrogen Launches With Faster Agentic AI, 1 Million-Token Context ... - Pulse 2.0 · pulse2.com

[17] Introducing Gemini 4 Hydrogen for Agentic Workflows - Google Blog · blog.google

[18] Anthropic Sees AI Agent Swarms Patching Flaws Within a Year - BankInfoSecurity · bankinfosecurity.com

[19] New Expel CEO Puts Agentic AI at Center of SOC Strategy - BankInfoSecurity · bankinfosecurity.com

[20] Flashpoint launches agentic AI to automate cybersecurity threat investigations - SC Media · scworld.com

[21] Red Hat Unfurls OpenShift Platform for the Agentic AI Era - Cloud Native Now · cloudnativenow.com

[22] ScienceLogic(r) Advances Agentic AI Investigation with New Specialist Agents in Skylar AI(tm) 2.8 · finance.yahoo.com

[23] Extreme Networks' big bet on agentic AI Ops is paying off - and there's much more to come · wifinowglobal.com

[24] Recent advances in AI are creating a 'paradigm shift' for robotics - Berkeley News · news.berkeley.edu

[25] Recent Advances in AI Are Creating a 'Paradigm Shift' for Robotics | Research UC Berkeley · vcresearch.berkeley.edu

[26] Can AI agents program industrial robots without training data? Berkeley says yes · therobotreport.com

[27] At CoRL, NVIDIA and Partners Reveal How Robots Rehearse Delicate Work in Simulation · blogs.nvidia.com

[28] Meet Alexa for Shopping, your personalized, agentic AI assistant on Amazon · aboutamazon.ca
[29] Apple Reportedly Wants Siri AI to Be More Agentic Next Year | PCMag · pcmag.com
[30] Apple explores agentic AI for Siri so it can finish tasks on its own in iOS 28 - Mac Daily News · macdailynews.com
[31] One UI 9.5 will bring Agentic AI capabilities for most stock apps - SamMobile · sammobile.com
[32] One UI 9.5 Could Bring Agentic AI Across Samsung Apps - SammyGuru · sammyguru.com
[33] Q&A: OpenTable hypes agentic AI - Expedite | Kristen Hawley · expedite.news
[34] What to Learn About Meta's Muse AI Before It Learns Too Much About You - CNET · cnet.com
[35] 5 AI agents you can actually use on Android in 2026 (and who should use them) · androidauthority.com
[36] Dialpad: Ada Combination To Build Agentic AI Customer Experience Platform Serving Over ... · pulse2.com
[37] Parloa Joins OpenAI Alpha Model Testing Program - CMS Wire · cmswire.com
[38] Decagon Missions: Can Purpose-Built Orchestration Win Enterprise CX? · futurumgroup.com
[39] Rescale Expands Agentic Digital Engineering with Customizable AI Agents for R&D Teams · prnewswire.com
[40] Siemens' HyperLynx brings physics-validated agentic AI to electronic systems verification · eejournal.com
[41] DuPont Announces Collaboration with Microsoft and NobleAI to Accelerate R&D Innovation with AI · dupont.com
[42] DuPont Partners with Microsoft and NobleAI on Agentic AI for R&D - Unite.AI · unite.ai
[43] Advancing Battery Modeling with COMSOL Multiphysics(r): GPU Acceleration and Agentic AI ... · volta.foundation
[44] From Chips to Skills, NVIDIA and Korea Expand Sovereign AI Collaboration | NVIDIA Blog · blogs.nvidia.com
[45] NVIDIA AI Day Seoul Highlights Korea's Advances in Sovereign, Agentic and Physical AI · blogs.nvidia.com
[46] Korea's AI Ecosystem Advances and Secures Agentic AI With NVIDIA · blogs.nvidia.com
[47] Nvidia Unveils 'Agentic AI Stack' -- "Token Usage Doubling Every Month, Inference Cost Is the Key" · finance.biggo.com
[48] Dell expands AI factory with NVIDIA for agentic AI - DataCenterNews Asia Pacific · datacenternews.asia
[49] Our newest TPUs: 8tx and 8i Max, built for the agentic era - Google Blog · blog.google
[50] Google Cloud SC26 Agentic HPC and AI Clusters - CASRAI · casrai.org
[51] Agentic AI Transforms Banking: Scaling Struggles - Funds Society · fundssociety.com
[52] From 'Whether' to 'How': The Emerging Agentic Playbook for Community Banks and Credit Unions · thefinancialbrand.com
[53] How RIT alumni are shaping new careers: Agentic AI engineer · rit.edu
[54] Agentic AI and financial crime: How banks can stay ahead - Tieto · tieto.com
[55] Optalitix brings agentic AI to insurance pricing - FinTech Global · fintech.global
[56] Deutsche Bank and Ant Digital Technologies Release White Paper on Agentic Payments in ... · prnewswire.com
[57] How Addy AI is bringing agentic AI into mortgage lending - HousingWire · housingwire.com
[58] Adobe Brings Agentic AI to the Content Supply Chain with CX Enterprise Coworker · news.adobe.com
[59] Adobe Brings Agentic AI to the Content Supply Chain with CX Enterprise Coworker | Nasdaq · nasdaq.com
[60] Adobe Bets on Agentic AI Across Creative Cloud - Fstoppers · fstoppers.com
[61] Adobe integrates Agentic CX Enterprise Coworker into AI platforms - FutureCIO · futurecio.tech
[62] Adobe is going all in on agentic AI - The Verge · theverge.com
[63] FBI FLASH on ARTEX: Agentic AI Hacking Frameworks Move From Pen Testing Tool to ... · oodaloop.com
[64] Amazon Threat Intelligence observes threat actors running cybercrime campaigns with ... - AWS · aws.amazon.com
[65] Agentic AI Attacks Are Here: What We're Seeing with ARTEX, and How to Protect Your Org · salesforce.com
[66] ARTEX and AI-Powered Red Teaming Tools: What Defenders Need to Know · unit42.paloaltonetworks.com
[67] Beyond ARTEX: 45 Agentic Attack Frameworks, What They Hunt For, and How to Find Them · blog.gitguardian.com
[68] Cybersecurity firms, overwhelmed by AI attacks, are turning away new clients · nbcnews.com
[69] India's CERT-In issues AI Security Auditing guidelines targeting model poisoning, prompt ... · industrialcyber.co
[70] NIST Draft Maps Security Threats Facing Agentic AI Systems - CDO Magazine · cdomagazine.tech
[71] Agentic security starts with visibility | Cisco Duo · duo.com
[72] Inside Falcon AI Gateway: One Control Point for AI Traffic - CrowdStrike · crowdstrike.com

[73] The Next Evolution of the Agentic SOC: Protecting Enterprise AI - Palo Alto Networks Blog · paloaltonetworks.com
[74] Agentic AI Security: How OX VibeSec Proves the Value of Risk That Never Happened · ox.security
[75] Delinea to showcase runtime controls for secure Agentic AI adoption at GITEX Global 2026 · intelligentciso.com
[76] AI cyber resilience and the trust gap - SiliconANGLE · siliconangle.com
[77] Where To Draw The Line On Agentic AI's Autonomy - Forbes · forbes.com
[78] AI Agents Behaving Badly: Four Security Rules For The Agentic AI Age - Forbes · forbes.com
[79] Governance Is The Foundation Agentic AI Can't Afford To Skip - Forbes · forbes.com
[80] AI Agents Are Testing the Limits of Zero-Trust Security - IEEE Spectrum · spectrum.ieee.org
[81] Google Cloud Agentic Architect Certification | Google Cloud Blog · cloud.google.com
[82] Live online certificate brings AI expertise to finance professionals | Cornell Chronicle · news.cornell.edu
[83] How the CIA Is Preparing Its Workforce for Agentic AI - ClearanceJobs · news.clearancejobs.com
[84] From Prototype to Production: Building Autonomous AI Agents with NVIDIA Agent Toolkit · nvidia.com
[85] Anthropic IPO to Test AI Industry's Financial Viability - BankInfoSecurity · bankinfosecurity.com
[86] Why tokenisation is about more than putting traditional assets on blockchain · top1000funds.com
[87] Publishers: Agentic AI Isn't About Pricing, But Whether You Get Any Bids At All · adexchanger.com
[88] AMD vs. Intel: Which Is the Better AI Chip Stock to Own for the Next 3 Years Following the ... · theglobeandmail.com
[89] Do AI agents have minds? A philosopher explains - The Conversation · theconversation.com
[90] FSBO.com Launches Native iOS and Android Apps, Bringing Its Agentic AI-Powered Real ... · prnewswire.com
[91] Cloudflare CEO Likens AI Agent Optimization To A New Version Of SEO · searchenginejournal.com
[92] Cedars-Sinai Software Helps Build, Monitor AI for Research · cedars-sinai.org
[93] FORGE Lab uses agentic tools to identify memory safety issues in Rust - Microsoft · microsoft.com
[94] Exclusive: Agentic AI Startup Lumbra Raises $25M Series A - Tectonic Defense · tectonicdefense.com
[95] 6sense Repositions as Partner for Agentic GTM - B2BMX INSIGHTS - Demand Gen Report · demandgenreport.com
[96] ABAP Platform Becoming an Enterprise-Grade Execution Backbone for Agentic AI · news.sap.com
[97] Dropbox Bets Its 700M Users on Agentic AI - Futurum Research · futurumgroup.com
[98] Denodo and AWS Partner to Accelerate Trusted Data Access for Agentic AI - BigDATAwire - HPCwire · hpcwire.com
[99] Pactum: Procurement's Execution Layer for the Agentic AI Age · procurementmag.com
[100] Fin Operator Routines: Proactive AI or Just Scheduled Tasks? - Futurum Research · futurumgroup.com
[101] Command Alkon Launches Command Intelligence, Bringing Agentic AI Into Command Cloud · finance.yahoo.com
[102] Immediately Apply Oracle Key Vault 21.16 Update | database · blogs.oracle.com
[103] LiteLLM now routes every request to the right OCI Generative AI model with a trained classifier · blogs.oracle.com
[104] Dawn of War IV' Is Coming to GeForce NOW - NVIDIA Blog · blogs.nvidia.com
[105] ServiceNow (NOW) Stock Could Be Pricey On Fresh Agentic AI Partnership News · ca.finance.yahoo.com
[106] Mars brings agentic AI-assisted operations into its new Hyderabad GCC - CRN Asia · crnasia.com
[107] How Constellation is building a governed data foundation on Databricks · databricks.com
[108] Post-Quantum TLS 1.3 Inspection: Closing the PQC Gap | Zscaler · zscaler.com