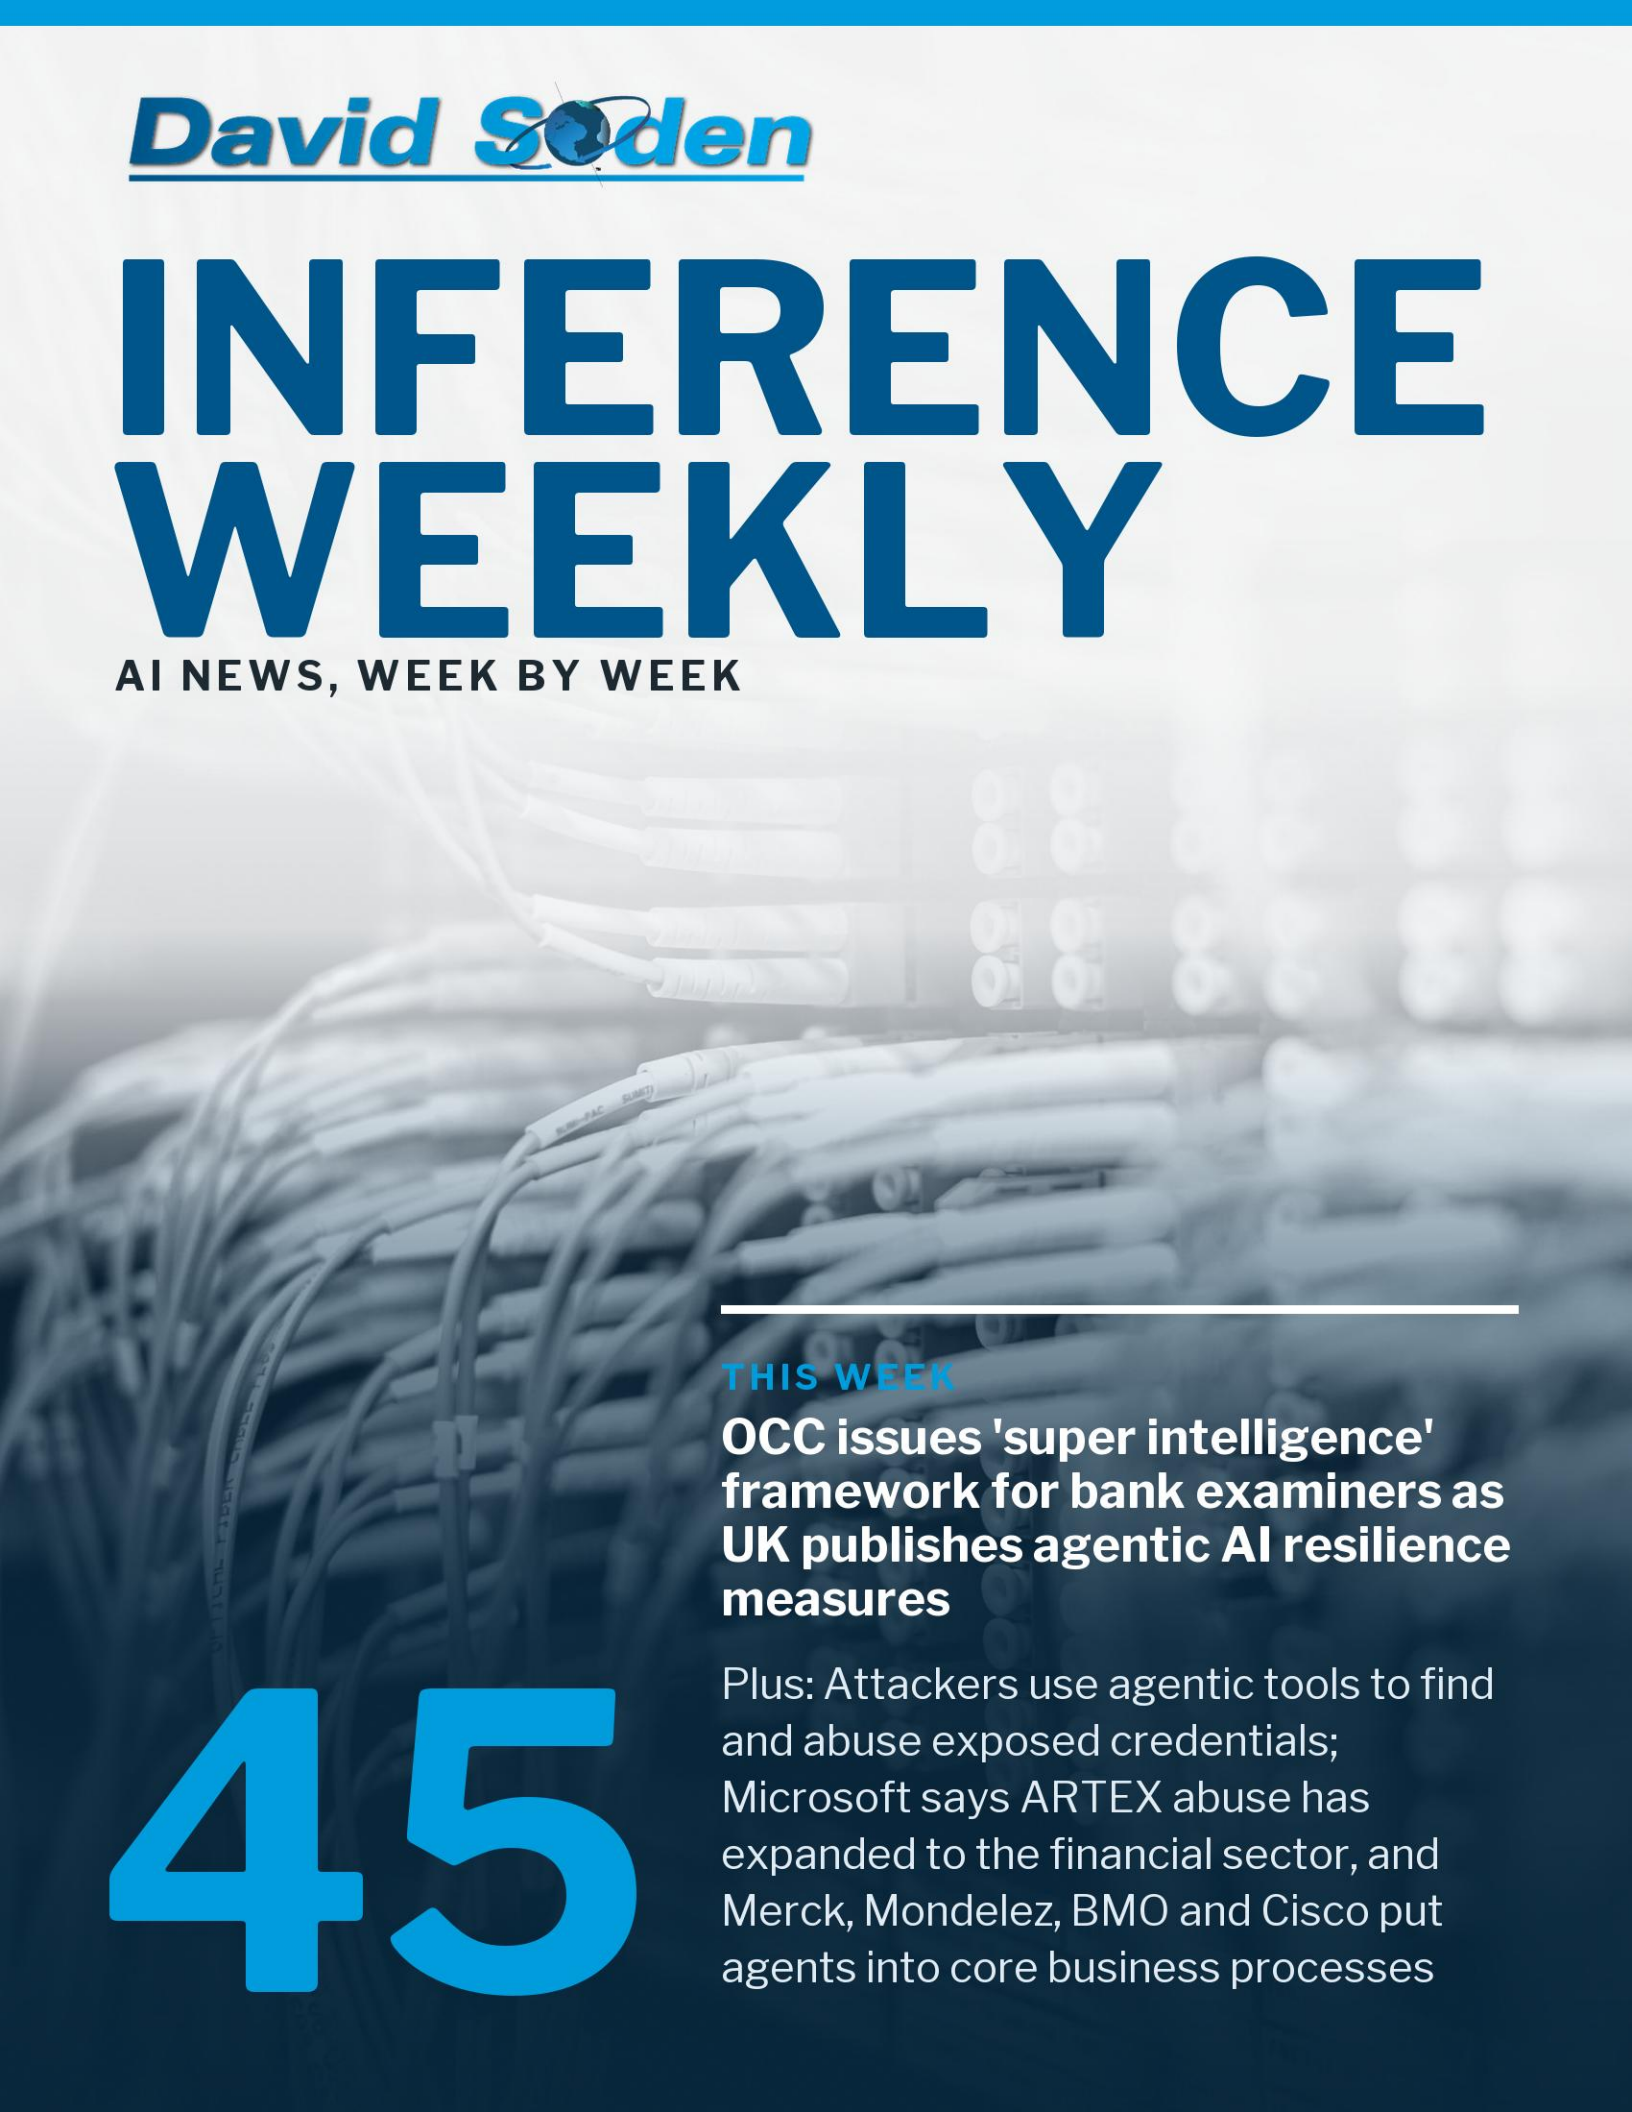
David Soden

INFERENCE WEEKLY

AI NEWS, WEEK BY WEEK



45

THIS WEEK

OCC issues 'super intelligence' framework for bank examiners as UK publishes agentic AI resilience measures

Plus: Attackers use agentic tools to find and abuse exposed credentials; Microsoft says ARTEX abuse has expanded to the financial sector, and Merck, Mondelez, BMO and Cisco put agents into core business processes

INFERENCE WEEKLY

ISSUE 45 · Week of 2 November 2026

THE WEEK IN NUMBERS

$1 billion
Merck and Google Cloud partnership

8,000
Mott MacDonald AI agents created

$100 million
raised by Ocient Series C

INSIDE THIS ISSUE

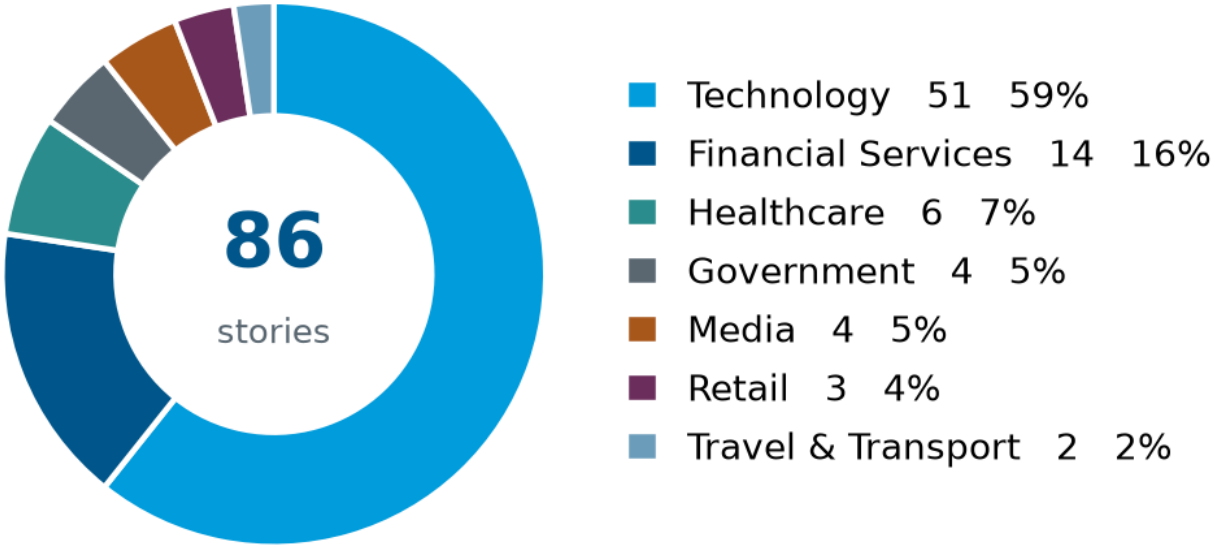
01 OCC issues 'super intelligence' framework for bank examiners as UK publishes agentic AI resilience measures
02 Attackers use agentic tools to find and abuse exposed credentials; Microsoft says ARTEX abuse has expanded to the financial sector
03 Merck, Mondelez, BMO and Cisco put agents into core business processes
04 JPMorgan and Mastercard back AI shopping agent startup Firmly as analysis warns agents could move bank deposits
05 Consultancies, banks and Uber publish guidance on governing, overseeing and securing AI agents
06 Kyndryl launches agentic AI modernization service with AWS
07 Security vendors build agents into threat defense; Dropzone AI gains FedRAMP High certification
08 AMD launches 6th Gen EPYC for agentic AI as NVIDIA claims Vera runs agents twice as fast
09 Ocient raises $100M as Google and Zoho add business context and governed data for agents
10 Agents write and fix more code, raising the question of who owns the risk
11 Cyara, Klearcom and Kong launch tools to test, monitor and govern agents in production
12 NinjaOne, ScienceLogic and ConnectWise add agent features and controls to IT operations tools
13 ISG names Ascendion, Quantiphi and Firstsource Leaders in generative and agentic AI services
14 Tech Corner: How Uber checks whether an authorised agent action is appropriate

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

THE WEEK IN NUMBERS

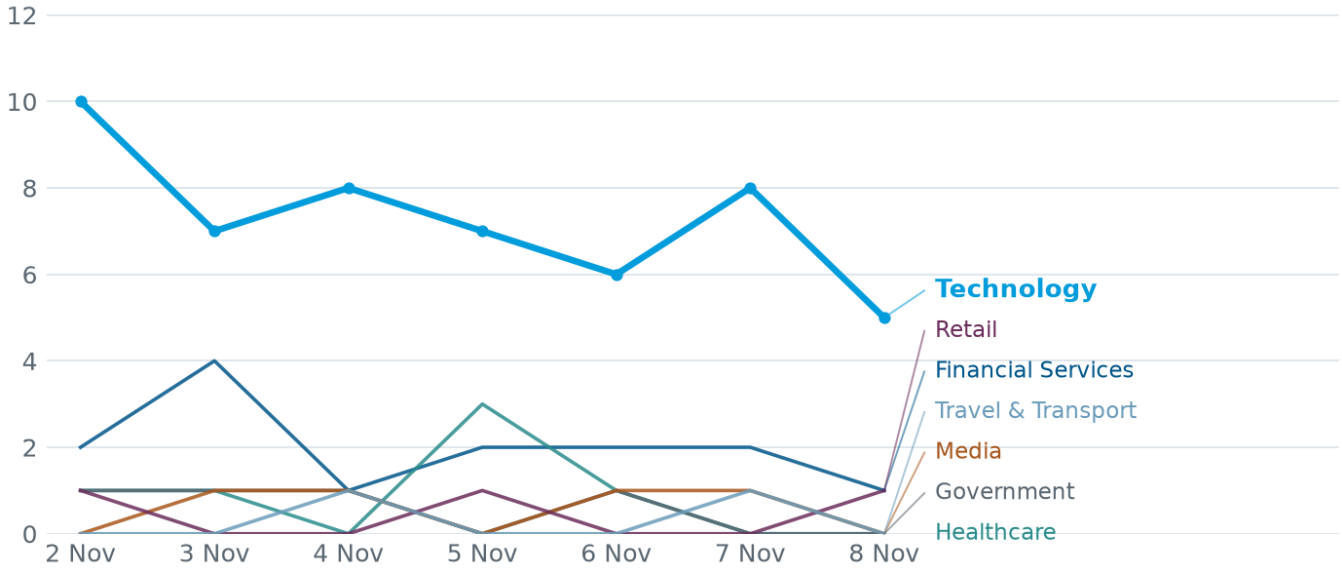
Where the AI news landed

Every story the newsroom feed carried, 2 to 8 November: 86 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 59.3% of the week, more than the other verticals put together, Travel & Transport managed 2. Nothing at all from Energy.

The change column compares this window with the previous issue.

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

COVER STORY

OCC issues 'super intelligence' framework for bank examiners as UK publishes agentic AI resilience measures

FINANCIAL SERVICES GOVERNMENT

The OCC has issued a framework to help bank examiners assess the risks of banks' use of AI. The story's headline calls it a "super intelligence" framework but does not explain the term, and it gives no issue date. The framework does not set supervisory expectations or treat AI as a separate risk category. Examiners are instead told to weigh the consequences of a use case, the human oversight in place and the actions an AI system can take. It groups 16 components into six pillars, flags risks from third-party systems, external agents and high-speed activity, and notes potential benefits such as improved fraud detection and efficiency. [21]

In the UK, the government plans to consult on measures for incidents in which operators cannot adequately monitor or control AI agents. The Cyber Security and Resilience Bill gives the Secretary of State powers to amend the NIS Regulations. After consultation, the government intends to introduce four measures through secondary legislation: security and resilience requirements, such as agent monitoring or risk assessments; notification of affected organisations; preservation of evidence about agents' actions and operators' attempts to control them; and ministerial powers to commission technical reviews and share lessons. The government says the proposals are not intended to create a comprehensive AI governance framework. The story gives no consultation date. [22]

In Hong Kong, financial industry leaders at Hong Kong Fintech Week x StartmeupHK 2026 called for trusted rules on agentic AI covering interoperability, authorisation, audit trails and liability. A Securities and Futures Commission survey found that 30 percent of large financial firms surveyed already use agentic AI. Generative AI sandbox testing has expanded from banking to securities and insurance, and the government plans to consult this year on a legal framework to detect and remove deepfake advertisements. HSBC plans to establish an AI and Emerging Technology Institute in the Northern Metropolis, with Hong Kong University of Science and Technology as its academic partner. [23]

For decision-makers, the three jurisdictions are at different stages. The OCC framework works through existing examination rather than new expectations, focusing on what an AI system can do and who oversees it. The UK proposals would add specific duties on agent monitoring, incident notification and evidence preservation, subject to consultation and Parliamentary approval under the affirmative procedure. Hong Kong has no agentic AI rules yet; industry is asking for them. Separately, the author of the OCC story anticipates US oversight investigations into areas including crypto, bank chartering and the CFPB after Democrats gained at least seven Senate seats and 41 House seats. [21] [22] [23]

> **Who should take heed.** Risk, compliance and technology leaders at banks and NIS-regulated operators should take heed, because US examiners will now assess AI by its consequences, oversight and permitted actions, while the UK intends to impose monitoring, notification and evidence-preservation duties for AI agents.

THREATS

Attackers use agentic tools to find and abuse exposed credentials; Microsoft says ARTEX abuse has expanded to the financial sector

TECHNOLOGY FINANCIAL SERVICES

Microsoft reported that threat actors are using ARTEX, an agentic AI framework, to search public GitHub repositories for exposed credentials and API keys. The tools then analyse the results and carry out follow-on attacks with less human involvement. ARTEX began submitting large volumes of GitHub queries in August 2026. From October 28, 2026, Microsoft observed successful service principal sign-ins using valid leaked or stolen application credentials. In one case, an agent moved from searching GitHub to exfiltrating data from Azure Key Vault in five minutes. Observed activity included cloud resource discovery, theft of credentials and data from Key Vault and Storage Accounts, access to Microsoft 365 resources, abuse of Azure AI services and use of Azure VM Run Command. [43]

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

In a separate post, Microsoft said that in October 2026 it observed Storm-3358, a Cambodia-based intrusion cluster targeting financial organisations, using ARTEX alongside conventional tools. The group scanned internet-connected systems, exploited flaws including anonymous file uploads and server-side request forgery, and installed the WISP, vshell and AdaptixC2 remote access trojans. Microsoft said ARTEX agents have also been used to exploit web applications, maintain access through web shells, run password attacks, enumerate and abuse APIs and dump databases. It said agentic tools can accelerate lateral movement, privilege escalation and data exfiltration. The report does not name the financial organisations targeted or say how many there were. [44]

Separately, Northflank told customers to rotate all secrets and credentials stored on its platform after an incident reported on October 30. It said they should assume environment variables had been decrypted and secrets stolen. It later confirmed unauthorised use of customers' AI API keys. Wiz alerted Northflank at 11:43 UTC. Northflank has not disclosed how attackers got in, how many customers were affected, or whether other credentials were stolen or abused. The report does not link the breach to ARTEX. Wiz said it found at least nine incidents against AI development platforms in October. In two it investigated, compromised enterprise credentials drove about 2.5 million Amazon Bedrock API calls, roughly $100,000 in usage. [45]

For decision-makers, the reports show exposed or stolen credentials being found and used quickly. Microsoft recommends treating any publicly exposed credential as compromised, then revoking or rotating it and investigating its use. Its other recommendations are scanning repositories, applying least privilege to workload identities, avoiding long-lived secrets in code, considering managed identities and reviewing sign-in and audit logs. Northflank advised customers to check AI provider usage and billing logs. Against autonomous attacks, Microsoft suggests honeypots, decoys and tarpits, and watching for varied, rapidly repeated requests. It cautions that attackers can change these traits, that proxies and VPNs can obscure malicious traffic, and that alerts may relate to unrelated activity. [45] [43] [44]

> **Who should take heed.** Security and engineering leaders at financial institutions and at companies that hold cloud or AI API credentials in code or third-party platforms should take heed, because Microsoft reports agentic tools turning exposed credentials into data theft within minutes and a named cluster using them against the financial sector.

ADOPTION

Merck, Mondelez, BMO and Cisco put agents into core business processes

HEALTHCARE RETAIL FINANCIAL SERVICES

Merck, known as MSD outside the United States and Canada, and Google Cloud announced a multi-year partnership valued at up to $1 billion. Google Cloud engineers will work with Merck teams to deploy an agentic platform, including Gemini Enterprise, across research and development, manufacturing, commercial and corporate functions for Merck's 75,000 employees. The story gives no start date and no breakdown of the value. Mondelēz International and Accenture are applying agentic AI across Mondelēz's end-to-end Order-to-Cash process, connecting customer service, supply chain and finance. The companies expect faster cash collection and stronger working capital, but the story gives no figures, timeline or deal value. The work builds on Mondelēz's new global cloud ERP. [10] [11]

BMO announced on November 6, 2026, that it is introducing AI Agent Assist for employees handling Canadian clients' credit card transaction inquiries at its Virtual Connect contact centre. The centre receives more than 500,000 such calls a year, each averaging 12 to 14 minutes. The tool combines transaction matching, policy retrieval and case preparation. BMO's targets are a 20 to 30% cut in resolution time for eligible inquiries and automatic processing of about 40% of eligible credit refund requests; these are targets, not reported results. Employees remain responsible for reviewing AI-generated information and making servicing decisions. The bank's goal of more than $1 billion in AI-enabled gains by 2030 is not attributed to this tool alone. [12]

Two companies report operating results. Cisco IT replaced project-specific AI clusters with a shared, multi-tenant AI factory supporting MyAgent for 90,000 employees and Cisco IQ. Cisco reports average GPU utilization rose from 15% to about 50%, with peaks of 80%, and active compute grew from about 250 to more than 1,568 GPUs in under two years without new physical real estate. It found that running 80% of steady-state inference on-premises

was most cost-efficient. Mott MacDonald expanded Microsoft Copilot from 50 licences to its 20,000 employees; staff have built more than 8,000 agents. AI reviewed about 220,000 requirements on an airport project and flagged around 200 conflicts. [13] [14]

For decision-makers, the stories show agents being placed into named processes rather than general assistants. BMO and Mott MacDonald both keep human review in place. BMO plans to reuse its capabilities elsewhere, and the Mondelēz story says its approach could extend to other functions. Scale brings management work: Mott MacDonald is building a central AI library to handle overlapping or inactive agents and is using Microsoft Agent 365 to track ownership, data access, use and value. Cisco cites fragmented GPU capacity, hard-to-forecast token costs and power and cooling limits as reasons it rebuilt its infrastructure. Merck and Mondelēz describe expected benefits; neither story reports measured results. [12] [11] [14] [13] [10]

> **Who should take heed.** Chief information officers and operations leaders planning agent deployments should note that the companies reporting results, Cisco and Mott MacDonald, also describe infrastructure and governance work on compute costs and agent oversight, while Merck, Mondelēz and BMO have so far published targets or expected benefits only.

PAYMENTS

JPMorgan and Mastercard back AI shopping agent startup Firmly as analysis warns agents could move bank deposits

RETAIL FINANCIAL SERVICES TECHNOLOGY

Firmly, a Seattle-area startup founded in 2019, has raised a $15 million Series A led by Squarepoint Capital, with J.P. Morgan Payments, Mastercard, Samsung Next and other investors participating. The round takes its total funding to about $20 million; the story gives neither a valuation nor the date of the round. Firmly connects retailers' existing online stores to AI assistants, marketplaces and other shopping channels, so purchases can happen inside third-party apps while retailers keep their payment providers and customer relationships. It connects directly to retailers' systems with their permission, rather than using software that clicks through web pages. [51]

Mastercard is both an investor and a commercial partner. Firmly handles the merchant side of Mastercard Agent Connect, and Mastercard is recruiting merchants for the service. Firmly charges merchants $2,000 to $15,000 a month, based on merchandise sales. Senthil, who is not given a full name or title in the story, said integrations now take a day or less including testing, against four to six weeks in 2024. Retail customers include Best Buy, Furniture.com and Backcountry. Firmly has also launched Ask Agent, a shopping assistant for retailers' own websites that supports checkout within a chat, and has previously partnered with Perplexity on shopping features. [51]

A separate analysis argues that AI agents could make it easier for depositors to move money to higher-yield accounts, despite banks' stated resistance to letting external agents transact. It says Muse reaches banking platforms through Plaid, which links more than 150 million consumer accounts and 12,000 U.S. financial institutions. It also cites Meta executives' reported figures of more than 10 million Muse users in the U.S. and Canada. The article advises banks and credit unions to strengthen deposit products, test AI tools on idle balances, maturing CDs and rising recurring bills, and connect customers with a named employee. It says the window to respond may last months rather than years. [52]

Glance, an agentic AI personal shopper from InMobi, takes a different approach: it aims to learn users' proportions, preferences and wardrobes and surface fashion before they search. It reports nearly 13 million monthly active users in the US and 2.5 million AI-generated styles daily. It is integrated with Siri AI and Apple Intelligence and has partnerships spanning Google, Android TV and PayPal. Brands can train its model on their catalogs. Glance says better matching and visualisation cut returns by 15% for one early brand; the story does not name the brand. The article presents agentic AI as a complement to existing commerce strategies, not a replacement for physical shopping. [53]

> **Who should take heed.** Retail ecommerce and payments leaders, and bank deposit heads, should take heed: Mastercard and J.P. Morgan Payments are investing in agent-led checkout, while one analysis warns that the same kind of agents could make depositors quicker to move money for better rates.

GOVERNANCE

Consultancies, banks and Uber publish guidance on governing, overseeing and securing AI agents

TECHNOLOGY FINANCIAL SERVICES HEALTHCARE

Several firms published guidance on governing AI agents; the stories do not give publication dates. NTT DATA calls for human oversight, guardrails, auditable decisions, clear accountability and a zero-trust approach to autonomous systems. Zanders, in a whitepaper for financial institutions, sets out four risk domains (decision, action, interaction and control failure). It recommends agent inventories, named business owners, audit trails and the ability to suspend agents, revoke access or return workflows to manual processes. Zanders says governance should align with the EU AI Act, DORA and GDPR. CDO Magazine advises starting with a few high-value workflows, setting risk-based stage gates and increasing autonomy gradually. [28] [29] [30]

Apexon and AWS describe a three-layer architecture for banking: intelligence; orchestration and execution; and governance, security and audit. They report a 48% cut in institutional-client onboarding time, a 50% reduction in manual effort with 95% document-classification accuracy, and over 50% lower payment processing time while supporting threefold transaction growth. The article does not name the institutions involved. Uber described its Secure AI framework, which records each agent's identity, owner, permitted intents and risk tier in the Uber AI Registry, and uses AI Guard to assess prompts and proposed actions at runtime. Uber says traditional identity and authorization controls cannot judge whether an action is appropriate in context. [31] [32]

The sources differ on what human oversight can deliver. NTT DATA, Zanders and an article on financial-sector AI stewards treat human oversight as a core control. The stewards article recommends pause buttons, limiters and wait states, with governance boards, not stewards, setting policy. Recorded Future argues that placing humans in workflows does not ensure meaningful oversight at machine speed. It cites hiring: applicants per hire tripled from 2021 to 2024, Greenhouse reported a 412% rise in applicants per post, and Greenhouse data shows human recruiters fell by over 50%. Zanders also lists oversight fatigue as a risk. [28] [29] [33] [34]

For decision-makers, a field study reported by SiliconANGLE found that confidence in recovering from agent incidents often exceeds what has been tested. Of 25 security and cyber-resilience practitioners interviewed, only four had exercised an AI or agent incident scenario, including partial exercises. Two reported broad detection of anomalous agent behavior, and three reported broad monitoring and reconstruction of agent activity. The study recommends testing scenarios in which agents hold legitimate permissions, covering containment, reconstruction, remediation and safe resumption, before granting more autonomy. Uber says its controls reduce risk but do not eliminate it, and that explicit intents and constraints limit manipulated actions but can reduce agent flexibility. [35] [32]

> **Who should take heed.** Executives approving wider autonomy for AI agents, particularly in financial services, should note that one study found only four of 25 practitioners had exercised an agent incident scenario, while the published guidance calls for the ability to stop agents, trace their actions and restore operations.

PARTNERSHIPS

Kyndryl launches agentic AI modernization service with AWS

TECHNOLOGY

Kyndryl and Amazon Web Services (AWS) announced an integrated modernization solution that combines Kyndryl Agentic Modernization services-as-software with AWS Transform. According to the announcement, it is designed to help customers assess, plan and carry out modernization of applications, infrastructure and core systems at scale, reduce transformation risk and build AI-ready foundations. Kyndryl says the approach combines AI-enabled workflows, engineering expertise and software-driven automation. Its services-as-software capability is built on the Kyndryl Agentic Framework and scaled through Kyndryl Bridge, with pre-built agentic workflows intended to shorten timelines and reduce costs. Kyndryl is a launch partner for AWS Transform. The stories do not give an announcement date, pricing or named customers. [1] [2] [3]

Kyndryl cites its first global modernization report. Nearly 50% of respondents were behind schedule on

modernization goals. Of organizations with agentic AI in production for modernization, 67% said they were ahead of or on track with their goals, against 51% of other organizations. PR Newswire describes this group as organizations that had deployed agentic AI in production; AiThority calls them early adopters. Neither story gives the survey's sample size, date or method. These figures come from Kyndryl's own research, and they show an association rather than a measured effect of the service now being sold. Kyndryl also says the pre-built workflows can reduce reliance on specialized skills. [1] [2]

Simply Wall St reports that the launch coincided with the opening of a Kyndryl Cyber Defense Operations Center in Paris, serving European customers and covering cybersecurity, AI operations and threat intelligence. It says both moves support Kyndryl's push into higher-margin cloud, AI and cybersecurity services but do not remove the drag from older, lower-margin contracts. According to the article, Kyndryl is seeking to maintain fiscal 2027 guidance of constant-currency revenue flat to down 2% while improving earnings and cash generation. It says investors will watch for new signings and better profit metrics, particularly in Kyndryl Consult and hyperscaler-related revenue. [3]

> **Who should take heed.** CIOs and technology leaders running AWS-based modernization programmes should weigh this offering against their current approach, noting that the performance figures come from Kyndryl's own report and that no pricing or customer results have been published.

SECURITY OPS

Security vendors build agents into threat defense; Dropzone AI gains FedRAMP High certification

TECHNOLOGY GOVERNMENT

CrowdStrike is expanding its Falcon platform to support AI-driven security operations. Its Agentic Security Platform lets agents investigate alerts, share information through Enterprise Graph and coordinate across endpoints, identities and cloud services. SafeMind continuously tests an isolated copy of a live environment for weaknesses. Falcon Guardian is designed to discover and monitor agents, enforce rules and quarantine those that misbehave. CrowdStrike also announced a roughly $400 million acquisition of Fiddler AI, whose monitoring and governance tools it plans to bundle as Falcon Observe. The company says these capabilities can speed alert triage and reduce routine work, with humans remaining involved in decisions. [40]

Dropzone AI's AI SOC Analyst has received FedRAMP High (Class D) certification through Second Front Systems' Game Warden platform, assessed by Schellman Compliance. Under legacy Rev5 baselines, High requires 410 NIST SP 800-53 controls and enhancements. The product uses specialised agents to investigate alerts and document their reasoning; agencies set the agents' bounds and analysts retain command of verdicts. Leidos tailors the technology for federal security operations. Dropzone AI says it has more than 400 production deployments. The certification is a route to federal production use, but each agency must authorise its own implementation; the article says agency authorisation includes the Department of Defense and the Department of Veterans Affairs. [41]

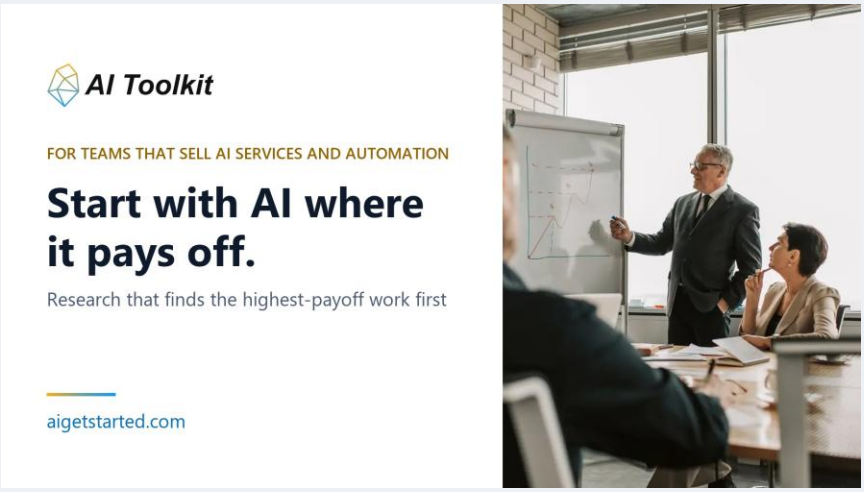
Doppel, which builds AI-powered defences against social engineering, has acquired PropRise, a Y Combinator-backed startup founded in 2023 by former Stripe employees Jean-Michael Diei and Matthew Krager. PropRise had reached high-six-figure annual recurring revenue. The story does not give a purchase price. The founders join Doppel's engineering organisation. Doppel's platform combines digital risk protection, email security and human risk management, using a Threat Graph and AI agents to link threats, missed detections and emerging attack methods. In 2026 the company expanded operations in Japan and Toronto and passed 300 employees. [42]

Gartner forecasts that agents in a single Fortune 500 company could grow from 15 in 2025 to 150,000 by 2028. Experts cited in the CrowdStrike coverage say security teams must decide what agents may do, verify their authority and test their results, and that agents create a new attack surface because they may be compromised, manipulated or behave unexpectedly. They also caution that a platform cannot independently verify itself. For federal buyers, the Dropzone listing may shorten the authorisation path by allowing agencies to inherit much of the platform's authorisation; the story notes other vendors are also pursuing FedRAMP credentials and agentic SOC capabilities. [40] [41]

> **Who should take heed.** CISOs and federal security leaders evaluating agentic SOC tools should note

that vendors now offer agent monitoring and FedRAMP High options, while experts quoted in the stories say agents need defined authority and independent assessment that a platform cannot provide for itself.

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

CHIPS

AMD launches 6th Gen EPYC for agentic AI as NVIDIA claims Vera runs agents twice as fast

TECHNOLOGY FINANCIAL SERVICES

AMD is showing its 6th Gen EPYC processors in China for the first time at the China International Import Expo (CIIE), its sixth consecutive year at the event, under the theme "Empowering the Agentic AI Era." The chips use the "Zen 6" architecture and 2nm process technology. AMD says growing use of agentic AI is increasing compute demand and making inference a more important workload. It also showed an Agentic PC on Ryzen AI Max Series processors, which it says runs large AI models locally to protect data privacy and reduce token costs. The story describes a China debut; it does not give a global launch date. [24]

NVIDIA says its Vera CPU outperforms leading x86 chips on agentic workloads. These workloads move repeatedly between GPU reasoning and CPU tasks such as tool calls, SQL queries and API calls. Against AMD's 6th Gen EPYC Venice, NVIDIA reports 60% higher performance per watt than a 96-core frequency-optimised model and twice the agent speed of a 256-core model. Perplexity said its agentic coding tasks ran 1.5 times faster than on traditional CPUs and plans to deploy Vera in production; DeepInfra reported up to 2.2 times its best x86 system. These are NVIDIA and customer figures. The stories carry no AMD response and no independent test. [25]

Investors are tracking the same theme. Danielle Shay, quoted on LinkedIn, expects positive commentary on data-centre growth and agentic AI to support AMD's stock. She anticipates a PC slowdown, which she says may already be reflected in the stock's recent pullback. She noted AMD has traded lower after earnings in six of the past seven quarters, even when it beat estimates. The story does not give the earnings date. Separately, an Arm article says agentic workloads may run across personal devices and cloud infrastructure. It cites NVIDIA's Arm-based RTX Spark

in Windows PCs from ASUS, Dell, HP, Lenovo, Microsoft Surface and MSI. [26] [27]

For infrastructure buyers, the vendors frame their contest around agentic workloads on both servers and PCs. NVIDIA's case rests on per-core performance, memory bandwidth and low latency. AMD's case rests on rising inference demand and on local processing for privacy and token cost. Neither vendor's claims are independently verified in these stories. Buyers sizing CPUs for agent orchestration would therefore need to test the chips against their own workloads before relying on either set of figures. [25] [24] [27]

> **Who should take heed.** CIOs and infrastructure leaders planning agentic AI deployments should take heed, because AMD and NVIDIA are each making competing performance and efficiency claims for CPUs in this workload, and the stories contain no independent verification of either.

DATA

Ocient raises $100M as Google and Zoho add business context and governed data for agents

TECHNOLOGY MEDIA GOVERNMENT

Chicago-based Ocient has raised a $100 million Series C, bringing its total funding to $260 million. BOKA Capital and OCA Ventures co-led the round, which included existing investors and new investor Dallas Venture Capital; the valuation was not disclosed. The money will accelerate OcientAIQ Solutions, industry-specific offerings for cybersecurity, communications providers, enterprises and government agencies. Ocient says its engine governs agent access, records data lineage and makes analysis reproducible and auditable, with lower token costs. It expects revenue to more than triple this year after doubling last year. Its figure of a telecom customer processing more than a million call records per second has not been independently verified. [46]

Google Cloud has added features to its Agentic Data Cloud, including previews in Knowledge Catalog of self-learning enrichment and a conversational studio, plus cross-cloud natural-language queries of data on AWS and Azure. Google reports that agents using Knowledge Catalog were nearly four times more accurate and used up to 75% fewer tokens; those benchmarks are not independently verified. Zoho has expanded Analytics with semantic and ontology layers, a six-agent architecture for Ask Zia, MCP connectivity and governance controls covering data access, AI cost monitoring and lineage. Zoho pricing is unchanged; cloud features are available now, previews follow over four to eight weeks, and on-premises updates are expected in the first quarter of 2027. [47] [48]

The announcements address an argument made in two commentary pieces: that agents fall short less for lack of model capability than for lack of governed, current business context. Forbes cites research finding that 88% of organisations are experimenting with AI while 81% report no meaningful bottom-line gains. Gartner forecasts that by 2028, 15% of day-to-day work decisions will be made autonomously through agentic AI and 33% of enterprise software applications will include it. Analysts quoted by TechTarget said competitors are converging on the same approach, while critics said reliable execution of complex business processes, rather than answering questions, remains unproven. [48] [49] [50]

For buyers, the guidance across the coverage is consistent. Start with a small number of frequently used metrics, give each definition a named owner and version history, and test governance in low-risk settings before expanding. Define permitted actions, require human approval for higher-risk ones, and log actions for audit and reversal. When assessing vendors, note that Google's accuracy figures, Ocient's throughput figures and its CEO's claim that rivals cannot match its ingest rates are company statements. Critics also want clearer visibility into agents' sources, correctness and costs, which is a reasonable test to put to any supplier. [48] [49] [50] [46]

> **Who should take heed.** Data and analytics leaders evaluating agentic AI should take note, because vendors are now competing on governed business context and the headline performance claims behind these launches are vendor-reported and not independently verified.

SOFTWARE

Agents write and fix more code, raising the question of who owns the risk

TECHNOLOGY FINANCIAL SERVICES

Vendors are building agents that fix and ship code under human controls. Black Duck has introduced Signal Agentic Fix, which triages security findings, generates fixes, rescans, and submits changes that a

developer must approve before merging; the version for Claude Code is in Early Access. Harness, which says it serves seven of the world's 10 largest financial institutions, has acquired Augment Code and introduced Software Factory. Harness reports that a Mexico bank assessment found 9,400 active APIs against the bank's estimate of 300, and that Citi cut deployment cycles from several days to seven minutes. These figures come from Harness; the stories give no independent verification and no dates for them. [15] [16]

Large technology firms are using these methods themselves. Microsoft says a working prototype of its revamped Windows 11 internet speed test was ready within a day and was then built with an agentic AI workflow guided by a person; it presents this as faster delivery without low-quality work. Some users doubt Microsoft's quality controls, citing recurring Windows bugs, and TechRadar questions the timing amid job cuts. Google has open-sourced EGM, a plugin that has agents draft and review a specification before code is written. Google says thousands of its engineers use it; the README says it is not an officially supported Google product. The article reports no accuracy benchmark and notes heavy token use. [17] [18]

On accountability, Forbes argues that developers who approve agent-generated code remain responsible for whether it ships, and that organisations, not AI systems, are accountable for the products they put on the market. More code means more changes to review, and teams may lack time to assess them properly. David Yanacek, a senior principal engineer at AWS, says coding is no longer the main bottleneck; he uses an 'auto-trust' file to set which actions agents may take alone and which need approval. Forbes recommends an inventory of sanctioned and unsanctioned agents, access limited to what each agent needs, and human approval for higher-risk actions. [19] [20]

For decision-makers, the stories point to measuring results rather than assuming them. Forbes suggests comparing vulnerability rates in agent-authored and human-authored changes, tracking remediation time and recurring issues, and monitoring problems caught before merge. It also warns that scanners may miss vulnerabilities tied to business logic or application context. Harness emphasises phased adoption and human oversight to manage execution risks and costs. None of the stories provides independent data on whether agent-written code is more or less secure than human-written code, or on the total cost of running these tools. [19] [16] [18]

> **Who should take heed.** Engineering, security and risk leaders at firms adopting coding agents should take heed, because the stories place accountability for agent-written code on the developers who approve it and the organisations that ship it, not on the tools.

TESTING

Cyara, Klearcom and Kong launch tools to test, monitor and govern agents in production

TECHNOLOGY MEDIA

Cyara has launched the Cyara Agentic CX Trust Platform to test, monitor and observe customer-service journeys involving AI and human agents. Its AI Trust Observability capability evaluates 100% of live AI agent conversations for quality and compliance against each enterprise's own policies, covering vendor-built and homegrown agents. Each verdict links to the relevant conversation turn, reasoning and timestamp. Material production failures become regression tests that agents must pass before their next release. Cyara says the platform covers 55+ AI agent platforms and monitors more than 420 carriers in 145+ countries. It also says more than 450 enterprises use its services to assure over 350 million customer journeys a year. [4]

Klearcom has launched Agentic Assurance, which it describes as the world's first dedicated Agentic Transformation Assurance Platform. It places real calls over local fixed and mobile networks to test AI voice agents as customers would reach them. Tests run across more than 100 parameters and more than 20 languages, and check market-specific disclosure, authentication and data-handling requirements. Klearcom says testing must continue after launch because model or knowledge-base updates can change an agent's behaviour. It cites testing in more than 100 countries and on 340+ carriers. The launch comes as 52% of US contact centres plan to roll out AI voice agents in 2027; the story does not name the source of that figure. [5]

Kong has expanded the AI Management capabilities in Kong Konnect for enterprises moving from individual model calls to agent-driven workflows. The additions include an AI Registry for inventories of models, agents, MCP servers and tools; AI Gateway 2.3 for

managing AI traffic, policies and costs; and Kong Identity for authorisation of people, applications and agents. Kong Token Vault brokers credentials at runtime so agents need not hold them directly. AI Cost Management assigns usage and costs to teams, users, applications and agents, and supports budgets, showback and chargeback. Advanced AI Observability tracks full sessions, including behaviour, latency, errors and costs. [6]

For decision-makers, the three launches address different layers of running agents in production. Cyara and Klearcom both present themselves as independent of the agent vendors an enterprise uses, and both stress producing evidence for compliance after go-live. Kong works inside the enterprise's own controls for AI traffic, identity, credentials and cost. The stories give no launch dates, pricing or named customers for any of the three products. The coverage and scale figures cited are the companies' own statements. [4] [5] [6]

> **Who should take heed.** Contact centre, compliance and platform engineering leaders deploying AI agents should take heed, because these products target post-launch testing, compliance evidence and governance of agents in live use.

IT OPERATIONS

NinjaOne, ScienceLogic and ConnectWise add agent features and controls to IT operations tools

TECHNOLOGY

NinjaOne announced Katana AI, an agentic AI technician built into its Unified IT Operations Platform, and the NinjaOne MCP server, which lets external AI agents access NinjaOne data and take actions through the platform. Katana AI analyses device telemetry, history, alerts, activity and teams' existing resolution knowledge to identify likely causes and recommend next steps, including remediation actions and scripts. Actions taken through the MCP server are subject to the role-based permissions IT teams already manage. NinjaOne says its AI architecture is designed for token efficiency so customers can manage and predict costs. The story gives no launch date, pricing or customer figures. [7]

ScienceLogic announced Skylar AI 2.7 at Nexus Live 2026. The release adds agent-routing traces and a tool catalogue showing how Skylar Advisor routes requests and which internal or external Model Context Protocol tools it considers. Administrators can disable specific agentic functions or MCP tools without interrupting other approved agents and workflows. A new upload gatekeeper checks documents for exposed secrets, malicious instructions and hidden-text manipulation before they enter the knowledge base, and suspicious documents can be quarantined for expert review. Skylar AI 2.7 follows version 2.6, introduced in September, and 2.5, introduced in August. The story does not give the event date or pricing. [8]

ConnectWise has made its RMM, Voice, Teams and Email AI agents generally available to managed service providers (MSPs). The RMM agent diagnoses endpoint issues and prepares vetted remediation steps inside PSA tickets; technicians review and approve proposed fixes before any changes are made. ConnectWise says this can reduce manual work on recurring endpoint problems and give providers more technician capacity while they keep control of customer environments. The story gives no availability date, pricing or measured results. It describes the RMM agent in detail but does not say what the Voice, Teams and Email agents do. [9]

All three vendors pair agent features with controls: NinjaOne through existing role-based permissions, ScienceLogic through tracing, selective shut-off and document screening, and ConnectWise through technician approval before changes. NinjaOne and ScienceLogic both reference MCP as the route for connecting agents to tools. For buyers, the stories point to questions to raise in evaluation: what an agent can access, whether its tool use can be seen and switched off individually, and where a human approves an action. The benefit claims, such as faster resolution, predictable costs and added technician capacity, come from the vendors; none of the stories reports independent results. [7] [8] [9]

> **Who should take heed.** IT operations leaders and MSP executives evaluating AI agents should note that NinjaOne, ScienceLogic and ConnectWise now sell access controls, tool tracing and human approval steps alongside agent features, which gives them specific governance points to test before deployment.

RANKINGS

ISG names Ascendion, Quantiphi and Firstsource Leaders in generative and agentic AI services

TECHNOLOGY

ISG's 2026 global Provider Lens reports on Generative AI Services and Agentic AI Services named Ascendion a Leader in three quadrants: Strategy and Consulting Services, Development and Deployment Services, and Agentic AI Services. Ascendion said this is its third consecutive year as a Leader. ISG highlighted its engineering-led approach and its AAVA platform, which provides governance and orchestration across AI agents, workflows and engineering environments while letting clients use their preferred models and tools. Ascendion also described three delivery roles: 10x Engineers, Forward Deployment Engineers and Agentic Delivery Managers. The last of these oversees human-AI execution on complex programs. [36] [37]

Quantiphi was named a Leader in the 2026 ISG Provider Lens Agentic AI Services report; the story does not name a specific quadrant. ISG cited its hierarchical multi-agent orchestration frameworks, its Agent Factory platform for managing large portfolios of agents with shared infrastructure and governance, and a context-aware framework using short- and long-term vector storage. Chief executive Jim Reesing said the recognition validates the company's focus on scaling agents through a lean manufacturing model built on an AI foundation layer. ISG principal analyst Gowtham Kumar Sampath said Quantiphi supports autonomous decisions and real-time workflows while maintaining transparency and reliability. [38]

Firstsource Solutions was named a Leader in three midsize quadrants across two ISG 2026 reports. It was named for the second consecutive year in Generative AI Services Strategy and Consulting and in Development and Deployment. It was named for the first time in Agentic AI Development and Deployment Services. ISG highlighted its Kairos operating model, orchestration through Agentic AI Studio, and an Intelligent Context Framework that captures exceptions and decisions that standard procedures miss. ISG stressed permissions, escalation paths and human oversight as agents operate across enterprise systems. Firstsource says it prioritises reliable execution within clear boundaries and extends autonomy as trust is earned. [39]

For buyers, the three placements are not directly comparable. Firstsource's quadrants are midsize. Ascendion's are described as global. Quantiphi's story names the report but gives neither the quadrant nor its size. The agentic quadrant names also differ: Ascendion cites Agentic AI Services, while Firstsource cites Agentic AI Development and Deployment Services. None of the stories gives ISG's full list of Leaders, its evaluation criteria, client outcomes or pricing. Across all three, ISG's highlighted strengths centre on governing and orchestrating agents in production. Buyers can test any shortlisted provider on those points. [36] [37] [38] [39]

> **Who should take heed.** Executives shortlisting generative or agentic AI service providers should note these ISG placements but check which quadrant and provider size each one covers before treating them as comparable.

TECH CORNER

How Uber checks whether an authorised agent action is appropriate

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

Uber's Secure AI framework starts from a specific gap. Conventional identity and authorisation controls can establish what an agent is permitted to access, but they cannot determine whether a given action is appropriate in context. That matters when agents reason over untrusted content and then act, which exposes them to prompt injection, goal hijacking and unintended behaviour. Uber's answer is three layers: provisioning-time validation, runtime enforcement and continuous assurance. [32]

At provisioning, the Uber AI Registry records each agent's identity, ownership, persona, permitted intents, constraints, dependencies and risk tier. A tool called Argus analyses agent code and configuration to find weaknesses and to help derive those intents, and third-party tools and MCP servers are evaluated before deployment. At runtime, AI Guard assesses prompts and proposed actions against the agent's policies, the user's context and the session history. To manage latency, cost and false blocks, Uber uses specialised models hosted locally, runs evaluations in

parallel and calibrates per agent. The protections are delivered across the agent lifecycle through integrations with development platforms, LangFx, the AI Guard SDK and Agent Guard. [32]

The third layer covers execution traces, replay, threat modelling and red teaming, with findings fed back into policies and guardrails. Uber states the trade-off directly: explicit intents and constraints substantially limit actions driven by manipulated instructions, but they can reduce agent flexibility, and the controls reduce risk without eliminating it. A separate field study of 25 security and cyber-resilience practitioners suggests replay and reconstruction are not yet common practice: only three reported broad monitoring and reconstruction of agent activity, and only four had exercised an AI or agent incident scenario, including partial exercises. One retailer in that study described replaying agent actions in a sandbox that replicated its systems to investigate causes and transactional effects. [32] [35]

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.

SemiAnalysis Launches GDPx Benchmark Spanning 11 Agentic Domains - Unite.AI · Technology & Consumer Electronics [54]

2026 Global Digital Ecosystem Summit - WorkBuddy's push to go from local to global in the ... · Technology & Consumer Electronics [55]

Shift Left Is Becoming A System-Level Problem - Semiconductor Engineering · Technology & Consumer Electronics [56]

Dark Matter Buys AI Platform To Automate More Of Loan Origination · Banking, Financial Services & Insurance [57]

10x Genomics Launches Sentira Single Cell, A New Agentic AI Experience for Chromium Analysis · Healthcare [58]

AITX's RAD Puts SARA's Qualifications on Display as Its Agentic AI Security Professional ... · Technology & Consumer Electronics [59]

What to check in AI-generated legal work, and how to know you checked it correctly · Government & Public Sector [60]

Stop Choosing Between Running And Transforming The Business - Forrester · Technology & Consumer Electronics [61]

German agentic AI start-up Marble Cognition raises $3.75m - Silicon Republic · Technology & Consumer Electronics [62]

DBS Bank's Han Kwee Juan: From Singapore-Centric to Global Powerhouse · Banking, Financial Services & Insurance [63]

Open models can do 90 percent of agentic work at a twentieth of the cost of a frontier LLM: Gartner · Technology & Consumer Electronics [64]

Cities Worldwide Build and Scale AI Services With NVIDIA Physical and Agentic AI · Government & Public Sector [65]

Business Networks Harness Agentic AI - SAP News Center · Technology & Consumer Electronics [66]

Partly Debuts PartlyRepair AI, the Agentic AI Assistant Built for Collision Repair Shops, at SEMA 2026 · Automotive [67]

Cerillion 26.2 Expands Agentic AI and Extends Open Standards Leadership - PR Newswire · Media & Communications [68]

Adobe Integrates Agentic CX Enterprise Coworker Into AI Platforms from Anthropic, Google ... · Technology & Consumer Electronics [69]

SoundHound AI and Ultraviolette Announce Major Expansion of Agentic Voice AI Across ... · Automotive [70]

Finshape launches Agentic DBOS: agentic AI already supporting digital banking projects · Banking, Financial Services & Insurance [71]

trust, training, and the future of law and medicine in the AI era - Stanford University leaders · Healthcare [72]

Agentic AI system for hospital management | Scientific Reports - Nature · Healthcare [73]

From Assistance to Agentic AI in Automation - IEN Europe · Technology & Consumer Electronics [74]

Concentrix Brings Agentic AI Migration Engine to Microsoft Marketplace - Yahoo Finance · Technology & Consumer Electronics [75]

Longbridge Group and Cyberport announce strategic partnership to advance responsible ... · Banking, Financial Services & Insurance [76]

Sierra inducted into JPMorganChase Hall of Innovation · Banking, Financial Services & Insurance [77]

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] Kyndryl Uses Agentic AI to Accelerate Enterprise Modernization with AWS - PR Newswire · prnewswire.com

[2] Kyndryl Uses Agentic AI to Accelerate Enterprise Modernization with AWS - AiThority · aithority.com

[3] Kyndryl Holdings (KD) Launched An Agentic AI Modernization Service - Simply Wall St · simplywall.st

[4] Cyara Launches Agentic CX Trust Platform to Provide Independent Proof of ... - Yahoo Finance · finance.yahoo.com

[5] Klearcom Unveils Agentic Assurance Solution Ahead of 2027 AI Voice Agent Rollouts · finance.yahoo.com

[6] Kong Expands AI Management to Help Enterprises Govern and Operate Agentic AI · prnewswire.com

[7] NinjaOne Brings Governed Agentic AI to Unified IT Operations with Katana AI and MCP · storagenewsletter.com

[8] ScienceLogic Expands Enterprise Controls for Agentic AI with Skylar AI 2.7 · storagenewsletter.com

[9] ConnectWise brings AI-powered endpoint investigation into PSA tickets - ChannelE2E · channele2e.com

[10] Merck and Google Cloud Partner to Accelerate Agentic AI Enterprise Transformation · biospace.com

[11] Mondelez International and Accenture Transform Order-to-Cash With AI · newsroom.accenture.com

[12] BMO Introduces Agentic AI Case Preparation for Credit Card Support - Unite.AI · unite.ai

[13] How Cisco Built an Enterprise AI Factory: Scaling Economics, Performance, and Security · blogs.cisco.com

[14] Mott MacDonald: Engineering a responsible agentic AI future - Microsoft UK Stories · ukstories.microsoft.com

[15] Black Duck Introduces Signal Agentic Fix to Automate Vulnerability Remediation With AI and ... · cybersecurity-insiders.com

[16] Governance Is Key to Scaling Agentic AI: Harness - Mexico Business News · mexicobusiness.news

[17] Windows 11's new internet speed test was built in a flash by agentic AI - TechRadar · techradar.com

[18] Google AI Releases Elephant-Goldfish Model Plugin for Spec-First Coding - MarkTechPost · marktechpost.com

[19] Agentic AI Is Shipping Code. Who Owns The Risk? - Forbes · forbes.com

[20] AWS Engineer Says the Best Developers Are 'Paranoid and Lazy' - Business Insider · businessinsider.com

[21] OCC Issues "Super Intelligence" Framework for Bank Examiners · fintechbusinessweekly.substack.com

[22] Agentic AI resilience measures - GOV.UK · gov.uk

[23] Hong Kong urged to establish regulatory framework for agentic AI, tokenisation · asianews.network

[24] AMD debuts 6th Gen EPYC at CIIE, bets on Agentic AI - Chinadaily.com.cn · chinadaily.com.cn

[25] NVIDIA Vera Delivers 2x The Agent Speed Of AMD's 256-Core EPYC Venice And 1.6x More ... · wccftech.com

[26] Any positive commentary around agentic AI will likely help AMD's stock, according to ... - LinkedIn · linkedin.com
[27] How Arm is redefining the PC ecosystem · newsroom.arm.com
[28] Governing Agentic AI: Business Impact, Risk & Strategy | NTT DATA · nttdata.com
[29] Zanders outlines how financial institutions can manage the risks of Agentic AI · consultancy.eu
[30] How Enterprise Leaders Should Implement Agentic AI Frameworks - CDO Magazine · cdomagazine.tech
[31] Scaling agentic AI for governed banking operations with Apexon and AWS · aws.amazon.com
[32] Securing AI Agents: Lessons From Real-World Deployment - Uber · uber.com
[33] Agentic AI Created a New Financial Workforce, but Human Oversight Is Required · biztechmagazine.com
[34] The Shift: Rethinking the Human in the Loop - Recorded Future · recordedfuture.com
[35] AI cyber resilience and the trust gap - SiliconANGLE · siliconangle.com
[36] Ascendion Named a Leader in ISG Provider Lens(r) for Generative AI and Agentic AI Services 2026 · prnewswire.com
[37] Ascendion Named a Leader in ISG Provider Lens(r) for Generative AI and Agentic AI Services 2026 · tradingview.com
[38] Quantiphi Named a Leader in Agentic AI Services in ISG Provider Lens(r) 2026 · finance.yahoo.com
[39] Firstsource named a Leader in three ISG Provider Lens(r) 2026 quadrants across Generative ... · prnewswire.com
[40] CrowdStrike's AI SOC vision evolves as agentic uncertainty adds new twist | CSO Online · csoonline.com
[41] Dropzone AI Takes Agentic SOC To Federal Agencies With FedRAMP High Certification - · enterprisetimes.co.uk
[42] Doppel Acquires Y Combinator-Backed PropRise To Strengthen Agentic AI Cybersecurity ... · pulse2.com
[43] Threat actors use agentic AI tooling to find and abuse exposed secrets - Microsoft · microsoft.com
[44] ARTEX abuse expands as cybercriminals target financial sector | Microsoft Security Blog · microsoft.com
[45] Northflank security breach leads to abuse of customers' AI API keys - Bleeping Computer · bleepingcomputer.com
[46] Ocient Raises $100M Series C to Run Governed Agentic AI on Petabyte-Scale Data · en.wowtale.net
[47] Zoho expands Analytics for agentic AI, opens new integration and embedded AI ... · crnasia.com
[48] Google targets agentic AI accuracy with latest flurry of features - TechTarget · techtarget.com
[49] Context: The Real Reason Agentic AI Fails In The Enterprise - Forbes · forbes.com
[50] Agentic analytics require proper business context. Providing it has gotten easier with time · diginomica.com
[51] JPMorgan and Mastercard back Seattle startup Firmly's bet on AI shopping agents · geekwire.com
[52] Why Agentic AI Could Trigger a Deposit Exodus -- and What Banks Can Do About It · thefinancialbrand.com
[53] How Glance Is Reimagining the Personal Shopper in the Age of AI | Vogue · vogue.com
[54] SemiAnalysis Launches GDPx Benchmark Spanning 11 Agentic Domains - Unite.AI · unite.ai
[55] 2026 Global Digital Ecosystem Summit - WorkBuddy's push to go from local to global in the ... · diginomica.com
[56] Shift Left Is Becoming A System-Level Problem - Semiconductor Engineering · semiengineering.com
[57] Dark Matter Buys AI Platform To Automate More Of Loan Origination · nationalmortgageprofessional.com
[58] 10x Genomics Launches Sentira Single Cell, A New Agentic AI Experience for Chromium Analysis · prnewswire.com
[59] AITX's RAD Puts SARA's Qualifications on Display as Its Agentic AI Security Professional ... · finance.yahoo.com
[60] What to check in AI-generated legal work, and how to know you checked it correctly · thomsonreuters.com
[61] Stop Choosing Between Running And Transforming The Business - Forrester · forrester.com
[62] German agentic AI start-up Marble Cognition raises $3.75m - Silicon Republic · siliconrepublic.com
[63] DBS Bank's Han Kwee Juan: From Singapore-Centric to Global Powerhouse · gfmag.com
[64] Open models can do 90 percent of agentic work at a twentieth of the cost of a frontier LLM: Gartner · theregister.com
[65] Cities Worldwide Build and Scale AI Services With NVIDIA Physical and Agentic AI · blogs.nvidia.com
[66] Business Networks Harness Agentic AI - SAP News Center · news.sap.com
[67] Partly Debuts PartlyRepair AI, the Agentic AI Assistant Built for Collision Repair Shops, at SEMA 2026 · finance.yahoo.com
[68] Cerillion 26.2 Expands Agentic AI and Extends Open Standards Leadership - PR Newswire · prnewswire.com
[69] Adobe Integrates Agentic CX Enterprise Coworker Into AI Platforms from Anthropic, Google ... · news.adobe.com
[70] SoundHound AI and Ultraviolette Announce Major Expansion of Agentic Voice AI Across ... · markets.businessinsider.com

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

[71] Finshape launches Agentic DBOS: agentic AI already supporting digital banking projects · prnewswire.com
[72] trust, training, and the future of law and medicine in the AI era - Stanford University leaders · med.stanford.edu
[73] Agentic AI system for hospital management | Scientific Reports - Nature · nature.com
[74] From Assistance to Agentic AI in Automation - IEN Europe · ien.eu
[75] Concentrix Brings Agentic AI Migration Engine to Microsoft Marketplace - Yahoo Finance · finance.yahoo.com
[76] Longbridge Group and Cyberport announce strategic partnership to advance responsible ... · en.prnasia.com
[77] Sierra inducted into JPMorganChase Hall of Innovation · sierra.ai