

David Soden

INFERENCE WEEKLY

AI NEWS, WEEK BY WEEK

THIS WEEK

AI Agents Ran A Full Ransomware Attack In Under Ten Hours, Then Left An 80-Page Report

Plus: Broadcom Adds Agent Identity, Security And Observability To Private AI Cloud, and CrowdStrike Extends OpenAI Partnership And Launches Falcon Guardian For AI Detection And Response

36

INFERENCE WEEKLY

ISSUE 36 · Week of 31 August 2026

THE WEEK IN NUMBERS

\$3.5 billion

NVIDIA convertible bond investment in MediaTek

\$550 million

Wonderful Series C funding

10 hours

agentic ransomware breach duration


INSIDE THIS ISSUE

- 01 [AI Agents Ran A Full Ransomware Attack In Under Ten Hours, Then Left An 80-Page Report](#)
- 02 [Broadcom Adds Agent Identity, Security And Observability To Private AI Cloud](#)
- 03 [CrowdStrike Extends OpenAI Partnership And Launches Falcon Guardian For AI Detection And Response](#)
- 04 [Palo Alto Networks Beats Quarterly Estimates And Pays \\$500M For Console](#)
- 05 [Surveys Find Agentic Adoption Running Ahead Of Governance, Security And Data Readiness](#)
- 06 [HPE, Arm, NVIDIA And SambaNova Tie Results And Launches To Agentic AI Demand](#)
- 07 [Anthropic Ships Claude Fable And Mythos 5.1 And Extends MCP To Hardware And Shopping](#)
- 08 [Beeline, GK Software, GovCore And DocuSign Open MCP Servers To AI Agents](#)
- 09 [Quiq, Genesys And Salesforce Push Agentic AI Into Contact Centres](#)
- 10 [Vendors And Engineers Turn To Audit Trails And Traceability For Agent Actions In Production](#)
- 11 [Google Releases EnvHarness And Mantis And Positions Gemini As An Agent](#)
- 12 [Google Bidding Change Exposes Shortfall In Agentic Advertising Returns](#)
- 13 [Genpact Launches Agentic Record-To-Report Suite With Outcome-Based Pricing](#)
- 14 [Huskeys, Acrab And Wonderful Raise Fresh Capital For Agentic AI And Agent Security](#)
- 15 [Tech Mahindra, Globant And Clearlake Build Agentic Practices With AWS And Google Cloud](#)
- 16 [Fobi AI Launches AgenticBrain Conversational Platform](#)
- 17 [Change Agents Completes Phase 2 Of Catch-Up Video Studio Platform With Caylent And AWS](#)
- 18 [Tech Corner: The Agent Data Path: Three Layers, One Missing Primitive](#)

THE WEEK IN NUMBERS


Where the AI news landed

Every story the newsroom feed carried, 31 August to 6 September: 117 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 65% of the week, more than the other verticals put together, Travel & Transport managed 2.

A change column appears once there are a few issues behind this one. This is number 3, so there is nothing honest to compare it with yet.

COVER STORY

AI Agents Ran A Full Ransomware Attack In Under Ten Hours, Then Left An 80-Page Report

TECHNOLOGY

Palo Alto Networks' Unit 42 documented an enterprise network breach completed in under ten hours, work the researchers say human red teams typically need about two weeks to finish. Two of the accounts date the disclosure to 2 September 2026; the victim is unnamed. Entry came through a public API. AI agents then ran reconnaissance, mapped the internal network, scraped code repositories for tokens, reached secrets management systems and took master admin credentials for root access. Specialised pivot agents validated access across cloud, identity, CI/CD, container and SaaS environments. On completion, an agent left the victim an 80-page security audit listing the vulnerabilities it had exploited. ^{[47] [48] [49]}

The accounts differ on who or what was driving. Two describe a human threat actor using frontier models and agentic attack frameworks; the third emphasises multiple purpose-built agents operating in parallel and adapting in real time rather than following scripts. One report counts over 50 MITRE ATT&CK techniques executed. One states no zero-day vulnerabilities were used, the speed and adaptation did the work. The attacker hijacked CI/CD workflows to steal cloud keys and repurposed the victim's own cloud AI services, masking activity as legitimate traffic. Multi-party code review blocked a Terraform backdoor attempt during the incident. ^{[47] [48] [49]}

Unit 42's guidance is containment at machine speed: revoke credentials, freeze CI/CD pipelines and isolate cloud accounts simultaneously rather than sequentially, using AI agents to automate the response. It also advises treating AI as core infrastructure, inventorying every model endpoint, API key, MCP gateway and tool integration, then applying rate limits and least-privilege policies, partly to avoid large token bills. Additional recommendations include behavioural detection tuned to AI patterns, governance and monitoring of AI infrastructure, and mandatory multi-party code review. One account

frames this as an escalation from the July 2026 JADEPUFFER attack, which used a single agent. ^{[47] [48]}

On the wider market, CrowdStrike's 2026 Threat Hunting Report says AI agent-triggered detections are rising 2.5 times faster than human-triggered ones, with one campaign sending nearly 200,000 model requests in two minutes. The same account cites 17,800 public agent add-ons, some impersonating Anthropic and OpenAI, and unpatched high-severity flaws including CVE-2026-82526 (SQL injection, CVSS 9.8) and CVE-2026-85620 (authorisation bypass, CVSS 9.2); it does not say these were used in the Unit 42 case. Funding and product responses named are HiddenLayer at \$100 million, AIR Security at \$50 million, Broadcom's AgentMinder and Okta's Agent SSO. ^[49]

Who should take heed. CISOs, heads of cloud platform and DevOps leaders should take heed, because the reported attack path ran through public APIs, code repositories, secrets stores and CI/CD pipelines they own, and the recommended countermeasure is simultaneous containment across all of them.

INFRASTRUCTURE

Broadcom Adds Agent Identity, Security And Observability To Private AI Cloud

TECHNOLOGY | FINANCIAL SERVICES | GOVERNMENT

Broadcom used VMware Explore 2026 in Las Vegas to announce a set of security, identity and observability products for agentic AI running on its Private AI Cloud. The line-up is AgentMinder, VMware vDefend and VMware Avi Load Balancer. Broadcom's stated argument is that agentic deployments widen the enterprise attack surface across agents, tools, datastores and large language models, in part because of newer protocols such as Model Context Protocol and agent-to-agent communication. The company positions Private AI Cloud as inference infrastructure with automated compliance guardrails intended to prevent data exfiltration. The stories do not give pricing for any of the three. ^{[54] [55] [56] [57]}

AgentMinder is a control plane that treats an agent as an enterprise identity tied to a declared mission, intent and approved tool set, authorising each action at runtime and logging sessions through OpenTelemetry. It supports existing authorisation systems via the AuthZEN standard and runs on-

premises, in a VPC or in public cloud. Broadcom says it runs the product internally on VMware vSphere Kubernetes Service and Google Cloud. The volume claims differ: one account says millions of customer and workforce API calls daily, another says up to 43 million daily. General availability is given as 31 August 2026, though other accounts describe it as unveiled at the event. ^{[58] [55] [57]}

vDefend extends Zero Trust lateral security with automatic discovery of agentic AI components, shadow AI monitoring for unauthorised use, and AI-generated intrusion detection and prevention signatures. Avi Load Balancer integrates with Kubernetes for AI-aware load balancing, web application and API protection, tool misuse prevention, zero-day detection and sensitive data protection. Separately, Broadcom unveiled the VMware AI Factory inside VMware Cloud Foundation, supporting over 150 models with Nvidia and AMD accelerators, and Tanzu Platform additions covering deny-by-default agent sandboxing, an MCP server marketplace, an Agent Buildpack and an AI Gateway logging prompts, tool usage and token consumption in Tanzu Hub. ^{[54] [55] [56] [57] [59]}

For buyers, the decision is whether to consolidate agent governance into the same stack that runs the infrastructure. The pitch is layered defence-in-depth with compliance-grade audit trails, aimed at high-risk workflows in finance, HR and IT. The counterweight is commercial: one account reports competitive pressure and some customer migration on pricing grounds since the acquisition, while noting many enterprises stay for maturity, stability and integration on on-premises AI. Note also that parts of the portfolio are not shipping, with AI gateways in VCF Private AI Services and Tanzu Platform described as future capability. Ask Broadcom for dates before committing. ^{[58] [57] [59]}

Who should take heed. CIOs, CISOs and infrastructure leads already standardised on VMware Cloud Foundation should take heed, because agent identity and runtime authorisation are now being sold as part of the platform rather than as a separate security purchase.

SECURITY

CrowdStrike Extends OpenAI Partnership And Launches

Falcon Guardian For AI Detection And Response

TECHNOLOGY | FINANCIAL SERVICES | GOVERNMENT

CrowdStrike announced an expanded partnership with OpenAI covering two directions of integration. OpenAI's Codex AI agents will be secured by CrowdStrike's Falcon Guardian, and OpenAI's GPT-5.6 Cyber reasoning model will be brought into the Falcon platform. Falcon Guardian controls Codex agent activity at runtime, providing live inventory, real-time visibility into agent activity, detection of compromised or unauthorised behaviour, and enforceable governance controls across endpoints, SaaS, cloud and browser environments. The stories do not give a date for general availability, pricing, or the number of customers involved. ^{[20] [21]}

The second leg of the deal places GPT-5.6 Cyber inside CrowdStrike's Frontier AI Readiness and Resilience (FAIRR) Service. According to the announcements, that combination pairs adversary intelligence, threat modelling, exploit validation and workflow orchestration to produce cyber risk assessments, attack path analyses and remediation priorities. One account specifies the model is for authorised defensive use cases and operates under expert supervision. Executives from both companies were cited on securing AI agents and using frontier AI for defence, but the stories name no individual from OpenAI and quote no financial terms. ^{[20] [21]}

Falcon Guardian was launched at CrowdStrike Fal.Con 2026 as an AI detection and response (AIDR) product, alongside Safe Mind. George Kurtz said AI is the latest battleground in cybersecurity and is accelerating the speed of attacks. President Michael Sentonas said visibility alone is insufficient once AI is embedded in an enterprise, and that Guardian gives runtime insight into what AI agents access and do so organisations can decide where to intervene. The product is positioned to find what CrowdStrike calls shadow AI: agents running across endpoints, who deployed them, and the risk they carry. ^[22]

For buyers, the decision point is whether agent activity is currently inventoried at all. The stories describe Guardian as extending CrowdStrike's endpoint detection and response approach to AI agents, using the same cloud-native, lightweight-agent Falcon architecture that also delivers automated protection, remediation and prioritised

vulnerability observability. Note the overlap: the OpenAI announcements and the Fal.Con launch describe the same product from different angles, and neither set of stories gives independent performance data, benchmarks or customer references.

CrowdStrike trades on NASDAQ as CRWD. ^{[20] [21] [22]}

Who should take heed. CISOs and heads of enterprise architecture running or piloting AI coding agents should take heed, because these announcements target agent inventory and runtime control, an area where the stories say visibility is currently incomplete.

EARNINGS

Palo Alto Networks Beats Quarterly Estimates And Pays \$500M For Console

TECHNOLOGY | FINANCIAL SERVICES | GOVERNMENT

Palo Alto Networks beat fiscal fourth-quarter estimates, reporting adjusted earnings per share of \$1.02 against an expected 98 cents and revenue of \$3.41 billion against a \$3.35 billion forecast. Revenue rose 34% from \$2.54 billion a year earlier, but the company posted a net loss of \$282 million, or 35 cents per share, versus net income of \$254 million, or 36 cents per share, in the prior year. Guidance was raised above analyst expectations: first-quarter revenue of \$3.30 billion to \$3.31 billion, and full-year revenue of \$14.10 billion to \$14.20 billion with adjusted EPS of \$4.16 to \$4.19. ^[18]

CEO Nikesh Arora attributed demand to rising AI-driven cyberattack risk, calling AI a "long-term tailwind." Shares have nearly doubled this year. The company said it held over 2,000 customer briefings following the launch of Anthropic's Mythos AI model, and cited breaches involving OpenAI and Hugging Face as evidence of concern over agentic AI attacks. CrowdStrike and Okta also gained on rising customer spending on cyber tools. The stories do not give the reporting date for the results, nor do they reconcile the 34% revenue growth with the swing to a net loss. ^[18]

Separately, TechCrunch reported Palo Alto paid \$500 million in cash and stock for Console, a two-year-old startup automating IT help desk tasks such as password resets, app access and troubleshooting. Terms were undisclosed when the deal was announced Tuesday. Console was founded in 2024 by Andrei Serban, raised \$29 million across a \$6.2 million

Thrive Capital-led seed and a \$23 million Series A coled by DST Global and Thrive, and was valued at \$157 million before the sale. Arora was an angel investor. The technology goes into the Cortex platform. ^[19]

The two accounts differ on Palo Alto's acquisition tally and pricing. TechCrunch calls Console the seventh acquisition of 2026 and puts Chronosphere at \$3.35 billion, alongside cyber startup Koi at \$400 million. The earnings report describes Chronosphere at nearly \$3.4 billion and adds CyberArk at \$25 billion, without a deal count. For buyers, the practical question is integration risk: a vendor absorbing seven companies in one year is consolidating IT service management, observability and identity into one stack, and Console's nearest rival, Serval, is now valued at \$1 billion. ^{[18] [19]}

Who should take heed. CISOs and IT operations leaders already on Palo Alto's Cortex or Prisma platforms should take heed, because seven acquisitions in a year, including Console, CyberArk and Chronosphere, will reshape the roadmap and the renewal conversation.

RESEARCH

Surveys Find Agentic Adoption Running Ahead Of Governance, Security And Data Readiness

TECHNOLOGY | FINANCIAL SERVICES | MEDIA

The adoption numbers in circulation do not agree with each other. A recent MIT survey cited by DevPro Journal puts AI agent adoption at 35%, with a further 44% planning to adopt. McKinsey's 2026 Global Survey on AI, cited around AgentsNexus India 2026 (Bengaluru, 4-5 September 2026), says 88% of respondents regularly use AI in at least one business function, but only about a third have scaled enterprise-wide. A KPMG UAE report puts UAE firms with AI agents in workflows at 97%, against what it describes as a global average of 87%. The stories do not reconcile the definitions behind these figures. ^{[60] [61] [62]}

Reported returns are narrower than reported adoption. DevPro Journal says only 16% to 30% of companies report productivity, time-to-market or customer experience gains, and only 31% to 45% report software quality improvements. Rubrik Zero Labs, surveying more than 1,600 IT and security leaders, found over 80% say AI agents require more manual oversight than the efficiency they provide.

Against that, GeekyAnts states its client work with Pillar Engine used agent-based architecture to cut manual effort by 99%, processing 10,000 pages with over 85% accuracy. That is a single vendor-supplied account, not survey data. ^{[60] [63] [61]}

Governance and control sit behind deployment. Rubrik reports 86% expect AI agents to outpace their security measures within a year, only 23% claim full visibility over agents (a number the report calls likely overstated), and 88% cannot roll back agent actions without disrupting systems; nearly half expect agentic systems to drive most attacks in the coming year. In manufacturing, nearly 75% expect autonomous AI deployment within two years while only 20% have mature governance. KPMG UAE reports 26% can enforce AI security strategies and 14% run modern integrated platforms. Solo.io cites 86% of enterprises lacking enforced AI identity access policies and has released agentdesktop under Apache 2.0. ^{[63] [64] [62] [65]}

Data readiness is the other constraint. Validity's survey of 500 global marketers found 25.6% rate CRM data 76-100% complete and accurate, 62% report direct revenue loss from data quality, and 75% name data integration and quality as the biggest challenge for agentic AI, yet only 32% prioritise data quality, unification and governance in AI investment. Some 44.7% let AI operate without human review. SDxCentral reports telecom subscriber data fragmented across legacy systems, and cites IDC forecasting 40% of Global 2000 job roles will involve working with AI agents by the end of this year. ^{[66] [67]}

Who should take heed. CIOs, CISOs and heads of data should take heed: the surveys consistently show agent deployment outrunning visibility, rollback capability and CRM or subscriber data quality, which are the controls that determine whether autonomy is recoverable when it goes wrong.

HARDWARE

HPE, Arm, NVIDIA And SambaNova Tie Results And Launches To Agentic AI Demand

TECHNOLOGY | GOVERNMENT | TRAVEL & TRANSPORT

Hewlett Packard Enterprise beat fiscal third-quarter analyst expectations and raised full-year guidance, citing record orders and backlog. Management lifted FY26 EPS guidance to \$3.75-\$3.85 with free cash flow of at least \$3.75 billion, and set FY27 targets of 13%-17% revenue growth, 14%-18% operating profit

growth, EPS of \$4.40-\$4.60 and free cash flow of at least \$5 billion. Networking revenue rose 10% to \$2.9 billion with orders up 36%; Cloud and AI revenue rose 25% to \$9 billion. After the quarter, HPE won a \$3.5 billion inference-server contract with a hyperscale client. Shares still fell 5.4% at Thursday's open. ^[50]

The gap between the numbers and the share reaction is the decision point. Bank of America Securities called the guidance conservative, arguing agentic AI and AMD Helios are excluded from networking projections. HPE itself flagged that the FY27 outlook reflects a shift toward AI revenue carrying lower gross margins, offset by operating expense control. Networks-for-AI orders reached \$700 million in the quarter, lifting the FY26 year-end target to \$2.5-\$3 billion, and AI systems orders totalled \$2.4 billion. Buyers negotiating AI infrastructure should note the vendor is trading margin for volume, and price accordingly. ^[50]

Arm reported record Q1 fiscal 2027 revenue of \$1.29 billion, up 22% year on year, split between \$715 million royalty and \$574 million licensing. Data-centre royalty revenue more than doubled. Non-GAAP operating income was \$531 million at a 41.2% margin, and Q2 guidance is roughly \$1.38 billion revenue and about \$0.47 adjusted EPS. The stock sold off recently on possible smartphone royalty declines and closed at \$242.59. At Arm's September Hot Chips event, four of six CPU research papers featured ecosystem partners: NVIDIA Vera, IBM's future Z/LinuxONE processor, Fujitsu MONAKA and the Arm AGI CPU. ^[51]

At Hot Chips 2026, SambaNova argued decode is now the inference bottleneck, at 75%-97% of model inference time across six frontier-model configurations, because agentic AI shifts work from single answers to ongoing reasoning and tool calling. Its SN50 RDU holds 432 MB on-chip SRAM to avoid repeated HBM access; scaling from 64 to 256 RDUs raised throughput from 200 to 500 tokens per second at 44%-51% model bandwidth utilisation. Separately, NVIDIA and MediaTek expanded their partnership around NVLink Fusion custom XPUs, with NVIDIA investing \$3.5 billion in MediaTek convertible bonds. No date was given for that agreement. ^{[52] [53]}

Who should take heed. Infrastructure and procurement leaders buying AI capacity should take heed, because vendors are repricing around inference and agentic workloads rather than training, which changes both cost structures and the margin profile of their suppliers.

MODELS

Anthropic Ships Claude Fable And Mythos 5.1 And Extends MCP To Hardware And Shopping

TECHNOLOGY | RETAIL | HEALTHCARE

Anthropic launched Claude Fable 5.1 and Claude Mythos 5.1, models aimed at coding, knowledge work and scientific research. Fable 5.1 is generally available at roughly 25% below the cost of Fable 5, driven by lower cache read pricing, with savings reported up to 45% on complex tasks. It ships with Enterprise Frontier Safeguards, offering zero data retention on customer-controlled cloud infrastructure, and reduced false positives in cybersecurity and biology. Mythos 5.1 runs more permissive safeguards and is restricted to vetted US organisations via Cyber Verification and Life Sciences Verification programmes, with expansion planned. No date for that expansion is given. ^[23]

On capability, Anthropic reports Fable 5.1 ahead of its predecessor on coding accuracy, problem-solving and long-duration tasks. Mythos 5.1 is credited with molecular design hit rates near 50%, a higher-resolution Venus elevation map built from NASA data, and speed-ups of up to 2.5 times across seven open-source deep learning models, cutting GPU costs. Safety testing places Mythos 5.1 as more capable but below higher risk tiers, with reinforced defences against misuse, prompt injection and distillation. EU AI Act compliance covers output watermarking and a detection API. The claims come from Anthropic's own testing and testimonials. ^[23]

Separately, Anthropic introduced the Model Hardware Standard as a research preview, letting AI agents discover and operate microscopes, liquid handlers, robotic arms and manufacturing equipment. Anthropic positions MHS as the physical-world counterpart to the Model Context Protocol, model-agnostic and reachable through MCP. The stated aim is cutting integration from weeks or months to hours or minutes. On IBM's Mixture of Experts podcast, IBM Principal Research Scientist Kaoutar El Maghraoui argued safety matters more here because physical errors differ from software ones. Simon Olson, founder of Toby, pointed to value locked in legacy industrial equipment. ^[24]

Anthropic also published blueprints for Claude-based shopping and merchant agents, covering catalogues,

carts, checkout, preference databases and purchase histories, with guardrails constraining pricing to real catalogue data and blocking manipulative upselling. Consumer appetite is unsettled: a Gartner survey found 11% willing to let AI make purchases, while an Accenture survey found 74% comfortable with routine tasks and 32% with partial purchase decisions. Advocacy groups raise surveillance and personalised pricing risks; experts warn merchants lack frameworks for disputed AI-generated purchases. Buyers evaluating agentic commerce should weigh the fraud and dispute gap before deployment. ^[25]

Who should take heed. Engineering, e-commerce and R&D leaders should take heed: cheaper frontier models, a hardware control layer and shopping-agent templates all landed at once, but the accountability frameworks for physical and purchasing actions have not.

INTEROPERABILITY

Beeline, GK Software, GovCore And Docusign Open MCP Servers To AI Agents

TECHNOLOGY | RETAIL | GOVERNMENT

Four vendors have announced Model Context Protocol servers within the same reporting window. Beeline launched Beeline MCP, which it describes as the first native MCP feature inside a vendor management system, giving AI agents a single governed connection to workforce data across sourcing, engagement and compliance. GK Software announced GK Agentic at NRF Europe in Paris, a retail AI framework and agent library that connects to retail backend systems via MCP. GovCore Inc., a Cardarelli Group company, made GovCore MCP generally available with 22 tools spanning licensing, renewals, continuing education, enforcement, payment reconciliation, reporting and forecasting. Docusign said it will open its MCP Server to all AI agents on September 30. ^{[37] [38] [39] [40]}

The common design point is that agents inherit human permissions rather than bypassing them. Beeline cites identity verification, role-based permissions and human oversight, and says high-stakes decisions on classification, pay equity and hiring stay human-driven. GovCore says agencies choose which models, tools and workflows are active, with audited tool calls, full audit trails and alignment to 20 NIST 800-53 control families, and that regulatory data stays in the US or

Canada. GK Agentic offers configurable human oversight. Docusign's server carries account-level controls, global infrastructure and multilingual support. None of the four stories provides independent verification of these controls. ^{[37] [38] [39] [40]}

Vendor independence is the second recurring claim. GovCore says building on the open MCP standard avoids vendor lock-in and supports government-approved cloud AI environments including Amazon Bedrock, Google Cloud Vertex AI and Microsoft Foundry. GK stresses platform-agnostic deployment across cloud, on-premise and hybrid to avoid reliance on a single model vendor. Docusign lists Claude, ChatGPT, Gemini, Copilot and Slack among the clients that will gain access, with its platform embedded in Oracle, Salesforce, Slack and Perplexity. Beeline positions its native integration against prior approaches that relied on multiple integrations or third-party wrappers. GK cites collaborations with FSAS Technologies and Fujitsu Research. ^{[37] [38] [39] [40]}

For buyers, the scale figures set the context. Beeline reports 25 years in market, over 450 customers, \$1 trillion in managed workforce spend and \$2.4 billion in customer savings; CEO Doug Leebby cited governance combined with workforce intelligence, and Chief Product Officer Colleen Tiner cited trust and human oversight. GovCore claims 99.9% uptime and hundreds of thousands of licensees. Docusign reports over 1.9 million customers and a Salesforce Partner of the Year award. GK became the largest global POS deployment provider in 2025. Beeline will demonstrate at the Staffing Industry Analysts' CWS Summit North America, dated September 28-29, 2026 in its release. ^{[37] [38] [39] [40]}

Who should take heed. CIOs, heads of procurement and compliance officers should take heed, because MCP is becoming the default connection layer between enterprise systems of record and third-party AI agents, and the permission, audit and data-residency terms are being set now.

CUSTOMER SERVICE

Quiq, Genesys And Salesforce Push Agentic AI Into Contact Centres

TECHNOLOGY | RETAIL | FINANCIAL SERVICES

Quiq announced Voice Assist, extending its agentic AI Assistants from chat, SMS and email into live voice calls. The stories do not give a launch date; Quiq says

the feature is available now. Voice Assist gives real-time adaptive guidance and can execute actions mid-call, such as sending a tracking link or gift voucher on the same conversation thread, without the agent leaving the call. It reuses Quiq's existing Process Guides, knowledge bases and verification guardrails. Quiq states that competitor tools may trigger backend workflows but do not send customer-facing messages during a call. CEO Mike Myer said the innovation is AI acting on suggestions live. ^[41]

Genesys expanded its Genesys Cloud Agentic Orchestration ecosystem with Adobe, AWS, Deepgram, ElevenLabs, Meta, Salesforce, ServiceNow and Sierra. Adobe covers bidirectional data flow between digital experience events and Genesys interactions; Salesforce and ServiceNow support Agent2Agent interoperability; Sierra and ElevenLabs agents run alongside Genesys Virtual Agents; Deepgram handles speech understanding; AWS supplies foundation model access; Meta extends WhatsApp messaging and voice. Genesys positions the platform as an orchestration layer that lets enterprises modernise without replacing existing systems. Separately, Salesforce was named a Leader in the 2026 IDC MarketScape for Worldwide Agentic CCaaS Platforms and Agentforce was ranked the number one Agentic AI Product in the 2026 G2 Best Software Awards. ^{[42] [43] [44]}

The stories disagree on market structure. One survey of seven agentic platforms lists Intercom Fin as a standalone option for businesses already using Intercom, alongside Cognigy, Sierra, Aissist.io, Zendesk AI, Fini and Decagon. A separate story states Salesforce acquired Fin, formerly Intercom, for \$3.6 billion. No date is given for the deal and the two accounts are not reconciled. Buyers evaluating a shortlist should confirm ownership before signing. Vendor-supplied performance figures also vary in basis: Aissist.io is credited with 83% automated resolution, Salesforce cites AI handling up to 80% of conversations, and Quiq reports a 20% reduction in agent conversation time. ^{[45] [46] [41] [43]}

CMS Wire argues the binding constraint is content infrastructure, not the model. Fragmented product information, policies and claims lead AI agents to present conflicting or outdated answers with confidence, and unlike humans they lack the judgement to weigh source reliability or escalate without explicit rules. The recommended pre-scale checks are a content audit naming authoritative sources per category, drift detection through

expiration dates and review triggers, risk-based escalation for legal, financial and safety topics rather than confidence thresholds alone, and feedback loops routing unresolved questions to content owners. Genesys frames governance controls as the enabler for regulated sectors. ^[46] ^[42]

Who should take heed. Customer service and CX leaders shortlisting agentic contact centre platforms should take heed, because vendor ownership is shifting mid-market and the stories indicate deployment success rests on knowledge governance the buyer must fix before scaling.

ADVERTISEMENT



**UNLOCK YOUR LEADERSHIP POTENTIAL:
WEEKLY LIVE PODCAST.**

 **Transformation**TM
— LEADERS PODCAST —

**TUNE IN THIS AND EVERY WEDNESDAY,
LIVE AT 12:00 NOON EST**

Join global business leaders for insights, strategies, and conversation.
Listen now at <https://www.transformationleaderspodcast.com>

Lead the Change, Don't Just Manage It.

Join transformation executives David Soden and Shawn Ennis every Wednesday at 12:00 PM EST for raw, candid conversations with industry giants. Gain the actionable digital strategies and leadership insights you need to successfully navigate complex business change.

www.transformationleaderspodcast.com

OPERATIONS

Vendors And Engineers Turn To Audit Trails And Traceability For Agent Actions In Production

TECHNOLOGY | FINANCIAL SERVICES | HEALTHCARE

Brian Kuan has released halo-record, an open-source Python package that writes each action an AI agent takes, tool calls, model calls, data accesses and approvals, to an append-only file in which every line carries a hash of the line before it. The chain makes edits or reordering detectable without vendor keys or permission, and the tool redacts secrets and personal data by hashing and pattern-matching. It plugs into OpenTelemetry spans and LangChain. The stories do not give a release date. The code is on GitHub, and Kuan plans to fund the work through hosted witness services. ^[32]

The limits are stated openly. The hash chain proves nothing was altered after writing; it cannot prove nothing is missing, because entries can be deleted before sealing. Kuan's answer is an independent witness service holding periodic counts and hashes as external proof of completeness. Current adopters are mostly engineers and security researchers, with pressure building from the AIUC-1 standard and insurance requirements. Help Net Security cites the July Hugging Face intrusion as evidence that auditing autonomous agent logs by hand is hard. The format is public and versioned, so implementations can verify each other. ^[32]

The same week, two engineering arguments landed on the same point from different directions. Itamar Syn-Hershko, founder of NeverBlink AI and BigData Boutique, argues databases cannot get a shadow mode: schema changes and migrations are irreversible in a way that pod restarts are not, and

general-purpose models risk functional hallucination, syntactically valid commands that are operationally harmful. He proposes a five-level autonomy ladder from observation to prohibited actions, plus machine-readable migration provenance. Dr. Aditya Vikram Kashyap, writing in India Today, names an agentic commit boundary, the moment a proposed action becomes a committed external change, and calls for a control plane handling identity, authority, retries, audit and recovery. ^{[33] [34]}

For anyone signing off a production deployment, the numbers matter. Kashyap cites benchmarks showing task success rates often under 50 percent and stability across repeat attempts below 25 percent, with multi-agent failures traced to coordination and verification rather than reasoning. KDnuggets reaches a similar conclusion across five enterprise cases, from SRE swarms to ERP reconciliation and security triage, warning that ReAct loops without rollback or idempotency produce ghost writes and data corruption. The Agentic Traceability and Testing Handbook, published 31 August 2026, covers the regulated end, mapping AI-assisted verification to DO-178C, ISO 26262, IEC 62304 and others via Visure and TASKING TBrn. ^{[34] [35] [36]}

Who should take heed. Engineering leaders and risk officers approving agent deployments should take note, because the week's material converges on the same claim: reliability and audit evidence now come from the runtime around the model, not the model itself.

PLATFORMS

Google Releases EnvHarness And Mantis And Positions Gemini As An Agent

TECHNOLOGY

Google Cloud AI Research, with Washington University in St. Louis and UNC Chapel Hill, released EnvHarness, a programmable layer that converts static agent benchmarks into environments that change as a policy trains. It wraps existing resettable environments through the standard reset() and step() interface without altering the simulator or verifier, using three composable components: Stage, Contract and Chain. An automatic designer, EnvRigger, diagnoses flaws from policy rollouts, writes wrapper components in Python and validates them on fresh

rollouts, revising up to five times per task. It is open source under Apache-2.0. No release date is given. ^[11]

Reported results across five benchmarks, ALFWorld, WebArena, SWE-bench Verified, OfficeQA and SpreadsheetBench, include a 9.0 point gain on out-of-distribution ALFWorld tasks and 9.8% fewer execution steps on SWE-bench Verified. The story states EnvHarness outperforms domain-specific environment generators and improves reinforcement learning results. The prerequisite is that environments must be resettable, which sets the practical limit on where this can be applied. For teams already running agent evaluation loops, the decision is narrow: the layer integrates into existing loops rather than replacing the underlying benchmark, so the cost of testing it is bounded. ^[11]

Separately, Google open-sourced Mantis, an agentic harness covering vulnerability identification, validation, reproduction and fixing. The stated problem is that conventional AI code scanning often returns true-positive rates below 7%. Mantis analyses repository history, past security fixes, architecture and threat models rather than brute-force file scanning, summarising files into a hierarchical tree that cuts token usage by 85%. It uses critic and reviewer agents, sandboxed reproduction for grounded evidence, and more than 15 tools communicating via shared disk storage. Google advises pairing lightweight models for classification with more powerful models for crash reproduction and fix generation. Code is on GitHub. ^[12]

On strategy, Google DeepMind SVP and Chief AI Architect Koray Kavukcuoglu said Gemini is moving from chatbot and language model to an agent that takes actions alongside humans, driven by advances in coding and software engineering. He said the underlying methods, deep learning, pre-training, reinforcement learning and optimisation, remain consistent; what changed is the operating environment, which requires models to infer intent, handle ambiguity and collaborate. He cited Gemini 3.5 for a deeper understanding of agentic workflows, and named Gmail, Sheets, Maps and Search as affected products. Asked what he wanted next, he said more intelligent models. ^[13]

Who should take heed. Engineering and security leaders running agent evaluation or AI-assisted code scanning should take heed, because both releases are open source and target measurable failure

modes: static benchmarks that stop teaching, and scanners whose true-positive rates fall below 7%.

ADVERTISING

Google Bidding Change Exposes Shortfall In Agentic Advertising Returns

MEDIA TECHNOLOGY RETAIL

Google Ads has changed how target-based bidding works for budget-limited campaigns, so algorithms now bid more aggressively to hit the set target, ROAS or CPA, regardless of prior overperformance. A campaign with a \$10 target CPA that had been delivering at \$5 will now drift towards \$10 unless the advertiser lowers the target. Efficiency previously banked becomes a spending threshold rather than a ceiling. The account of this is that Google earns more when advertisers spend more, and that Meta and TikTok face the same tension between advertiser return, publisher yield and shareholder growth. No date is given for the change. ^[26]

The commercial pressure behind that is set out plainly: Meta's most recent quarter showed rising revenue and rising costs driven by heavy AI investment, with advertising the primary means of funding it. The argument put is that platforms control the bidding rules, the signals, the inventory and the metrics, so autonomous buying systems cannot be assumed to be aligned with the buyer. The remedy proposed is independent baselines held outside platform reporting, media tied back to business economics, and explicit governance over automated buying. No specific figures for the Meta quarter are given. ^[26]

India offers a measured read on what agentic buying actually returns. Indian AdTech startups promised, a year ago, autonomous planning, buying, optimisation and reporting with minimal human input, and raised significant investment on it. After real-world trials, results are mixed: efficiency gains in mid-funnel work such as bid management, pacing and budget reallocation, producing modest CPA improvements largely by cutting human lag, but no genuine strategic intelligence. The most credible tests showed 14% incremental improvement, well below pitch estimates. AI often optimised narrowly for the assigned metric, sometimes at the expense of brand or wider objectives. ^[27]

The pattern now settling in India is supervised autonomy: humans retained for high-impact decisions, AI left with repetitive low-risk tasks, mirroring the earlier arc of programmatic. Performance held up in high-transaction, data-rich categories, quick commerce, subscription apps, fintech, and was weaker in considered-purchase sectors such as automobiles and real estate, and in brand campaigns needing long-term measurement. Vendors emphasising incrementality testing, transparent reporting and outcome-linked pricing are expected to survive client scrutiny. Separately, a Sciences Po seminar on 11 September 2026 with Koray Caliskan and Donald MacKenzie examines agent-initiated commerce including Amazon's Buy for Me, Walmart/OpenAI Instant Checkout and Visa Intelligent Commerce. ^{[27] [28]}

Who should take heed. Marketing and finance leaders running automated media budgets should take heed, because platform-set bidding rules can quietly convert their earned efficiency into higher spend unless they hold independent baselines and run their own incrementality tests.

FINANCE

Genpact Launches Agentic Record-To-Report Suite With Outcome-Based Pricing

FINANCIAL SERVICES TECHNOLOGY MEDIA

Genpact has launched the Genpact Record-to-Report (R2R) Suite, an agentic AI product for the financial close. It has three modular components: Journal Entry, Reconciliation and Intercompany, covering validation, anomaly detection, transaction matching, root cause classification and balance reconciliation before close. Genpact claims up to a 40% reduction in peak close effort, first-pass reconciliation yields above 95% and resolution of up to 99% of intercompany breaks in real time. Tenneco, a global manufacturer, is piloting the suite. The stories give no launch date, no pricing figures for the suite itself and no independent verification of the performance numbers. ^[8]

The pricing model is the part worth attention. Genpact sells outcomes rather than seats or tokens, taking the efficiency risk onto its own balance sheet: one cited example lowers a client's cost per invoice in accounts payable from \$5 to \$2. The company forecasts over \$1 billion in contracts by 2026 on this basis, with client Wesco processing 40% of invoices without human

intervention. The reporting describes outcome-based pricing as emerging but still rare at scale in enterprise software, so Genpact's position here is early rather than typical. ^[9]

For buyers, the controls matter as much as the savings. Genpact describes auditing, human overrides and transparent decision trails covering both AI and human steps, which is the condition for use in regulated finance. Its own guidance recommends diagnosing the finance value chain to find the 20% of workflows consuming 80% of manual effort, starting with accounts payable and invoice-to-cash, and standing up a centralised agentic finance command centre for governance, KPI monitoring and exception management. Responsible AI principles are aligned to NIST AI RMF, OWASP Top 10 for LLMs and DASf 3.0. ^{[9] [10]}

Two caveats sit in the material. First, AI contextual awareness remains limited and humans still handle exceptions, so the 40% and 99% figures describe assisted, not autonomous, close processes. Second, Genpact improves its agents using anonymised learning across clients, a data arrangement finance and legal teams should examine before signing. Role redesign is presented as essential, with finance staff moving from processors to supervisors and exception handlers, supported by reskilling. Decision-makers should treat the vendor's outcome guarantee as the negotiable item, and demand that the claimed baselines be measured against their own current close. ^{[8] [9] [10]}

Who should take heed. CFOs, financial controllers and shared-services heads should take heed, because the outcome-based contract shifts efficiency risk to the vendor but requires them to agree a defensible baseline and to accept cross-client anonymised learning on their process data.

FUNDING

Huskeys, Acrab And Wonderful Raise Fresh Capital For Agentic AI And Agent Security

TECHNOLOGY | FINANCIAL SERVICES | HEALTHCARE

Three companies working on agentic AI disclosed new funding. Huskeys, an Israeli cybersecurity firm, raised \$27 million led by Blackstone, taking total funding to \$35 million. Participants included Skinos Ventures, Shlomo Kramer's investment fund Yishay Yovel, funds from Okta and Zscaler, Merlin Ventures, executives

from Google, Microsoft, Palo Alto Networks, Cloudflare, Check Point, AWS and Intel, and private investors including actor Assi Cohen and mentalist Lior Suchard, who also advises the company. The Okta and Zscaler investments carry strategic partnerships covering integration of security solutions. No date is given for the round. ^[29]

Acrab raised US\$130 million from existing investors Vertex Ventures SEA & India, Vertex Growth, and institutional investors in Europe and Southeast Asia. The company recently left stealth with more than US\$350 million raised cumulatively and launched its GELIX 1-powered Agent Box platform, a full-stack AI computing platform spanning custom silicon, edge AI and software orchestration. It says the money will scale products, expand the ecosystem and build a next-generation platform, with industrial deployments and revenue targeted within the year. The story does not name a lead investor or give a closing date. ^[30]

Wonderful raised \$550 million in a Series C led by Insight Partners, at a \$5 billion valuation, double the \$2 billion it carried six months earlier. Founded in early 2025, the Israeli-Dutch company sells a model-agnostic AI operating system for agents, workflows and AI-native applications, and employs 650 staff across 35 markets. Salesforce invested for the first time as a strategic investor. Its differentiator is a forward-deployed engineer model, where senior technical staff embed in customer environments and hand control back over time, aimed at banking, telecommunications, healthcare, utilities, insurance and retail. ^[31]

For buyers, the three rounds point at different layers of the same shift. Huskeys argues the web application firewall layer, developed in Israel nearly 30 years ago and, it says, largely unchanged since, needs an agentic AI layer to handle autonomous agents; CEO Itai Gafni said such agents will dominate the internet and require new identity infrastructure. Huskeys names TikTok, Legoland, Ro, Blackstone and Hugging Face as customers. Wonderful's story notes the scalability of its engineer-heavy model remains uncertain and that agentic deployments still need human intervention. Acrab has not yet reported revenue. ^{[29] [30] [31]}

Who should take heed. CISOs and enterprise architects in regulated sectors should take heed, because vendors are now raising at scale on the premise that autonomous agents need new security, identity and integration layers that current infrastructure does not provide.

SERVICES

Tech Mahindra, Globant And Clearlake Build Agentic Practices With AWS And Google Cloud

TECHNOLOGY | FINANCIAL SERVICES | HEALTHCARE

Tech Mahindra launched an AWS Agentic Process Transformation Centre of Excellence, combining its Business Process Services process expertise with AWS cloud and agentic AI capabilities. The CoE covers telecom, healthcare, banking and financial services, retail and manufacturing. Its first joint solution, Collections Guru, an agentic collections agent, was deployed by Target Group in the UK and delivered roughly 40% efficiency gains in arrears management by autonomously optimising collection strategies. Tech Mahindra's President of Business Process Services framed the CoE as moving customers from experimentation to measurable impact with governance; the story does not name that executive or give a launch date. ^[14]

Separately, Hoxton Wealth partnered with AWS under its Forward Deployed Engineering initiative, backed by a \$1 billion investment that embeds AWS engineers with customer teams to co-develop and deploy agentic AI within days. Hoxton CTO Mansel Oliver cited AI-generated client reports to lift adviser productivity and an internal knowledge base to democratise expertise currently held by senior staff. Mark Jopling, an AWS director, said the aim is to close the gap between AI intent and deployment and leave Hoxton self-sufficient. No contract value, timeline or measured result is given. ^[15]

On the buy side, Clearlake Capital, a private equity firm with \$185 billion in assets, is giving portfolio companies access to Google Cloud's full AI stack: AI Hypercomputer with TPU accelerators, NVIDIA GPUs, data analytics, storage, multicloud cybersecurity and Gemini Enterprise for autonomous agents, with open-source and third-party models via Vertex AI. The models feed Clearlake's O.P.S. operational framework. Alteryx is already using the stack. This follows an OpenAI partnership supplying GPT-5.6 and ChatGPT Work to over 50 portfolio companies, plus earlier tie-ups with Databricks and West Monroe. Clearlake is buying from multiple vendors, not consolidating. ^[16]

Globant introduced the MuleSoft AI Pod, a delivery model pairing specialised AI agents with human

experts for API-led connectivity and integration, accessible through its Glob.AI platform. Globant claims migration timelines cut from months to weeks, up to 80% of manual integration tasks automated, delivery speed up 15-25% and a 35% improvement in design and development efficiency. Pods handle six to 24 integration flows monthly by size. Pricing is tied to outcomes rather than hours or seats, and access is via waitlist. Those figures are Globant's own; no independent customer result is cited, unlike Tech Mahindra's Target Group deployment. ^{[17] [14]}

Who should take heed. Operations and technology leaders in financial services should note that agentic AI is now being sold as a packaged delivery model with outcome-linked pricing, but only one of these four announcements, Tech Mahindra's Collections Guru at Target Group, carries a named customer result.

LAUNCHES

Fobi AI Launches AgenticBrain Conversational Platform

TECHNOLOGY | FINANCIAL SERVICES | HEALTHCARE

Fobi AI Inc. has launched AgenticBrain, a conversational platform that connects to a company's existing enterprise systems and lets customers search, book, modify orders, authenticate, track and request services through messaging channels including WhatsApp and iMessage. The stories state it makes existing systems conversational rather than replacing them, and that it carries persistent memory retaining customer history and preferences to personalise later interactions. The announcement stories do not give a launch date, pricing, named customers, or any named Fobi executive; one story refers only to unnamed "Fobi's leadership" describing AgenticBrain as completing the Flywheel. ^{[4] [5] [6] [7]}

AgenticBrain is positioned as the fourth component of what Fobi calls the Fobi Flywheel, joining FORTRESS (described as sovereign enterprise AI and intelligence), FIXYR (intelligent communication and orchestration) and AltID 3.0 (sovereign identity, continuous authentication and quantum secured trust). Fobi says the combined ecosystem spans AI, orchestration, identity, cybersecurity, quantum trust and customer interaction, and that the integration approach gives enterprises and governments greater control over proprietary data and AI infrastructure. The stories present this data sovereignty argument as Fobi's own positioning; none contains independent

verification, benchmark results or third party assessment. ^{[4] [5] [6] [7]}

Named target sectors are travel, hospitality, retail, financial services, telecommunications, healthcare, utilities and government, with one story adding ticketing. Fobi cites recurring revenue routes through licensing, deployments, usage, integrations and transaction-enabled services, and says the platform provides an interface where autonomous AI agents can act, authenticate and transact on a customer's behalf. For executives, the decision point is narrow: this is a launch announcement with no disclosed deployments, revenue figures or timelines. Buyers in the listed sectors should treat the sovereignty and agent-to-agent claims as vendor statements pending references and technical due diligence. ^{[4] [5] [6] [7]}

Who should take heed. Customer experience and digital channel owners in travel, retail, financial services, telecoms, healthcare, utilities and government should note the launch, but with no disclosed customers, pricing or dates in the announcement, any evaluation rests on independently verifying Fobi's sovereignty and agentic claims.

LAUNCHES

Change Agents Completes Phase 2 Of Catch-Up Video Studio Platform With Caylent And AWS

MEDIA TECHNOLOGY

Change Agents Corporation (Nasdaq: CHGA) announced the completion of Phase 2 development of its Catch-Up agentic AI video studio platform. Development was carried out with Caylent, Inc., described in the stories as an Amazon Web Services Premier Tier Services Partner. Story 25 states the work was led by Caylent and funded by an AWS grant; stories 16 and 18 describe the platform as created with or in collaboration with Caylent and do not mention grant funding. The platform is intended to generate personalized short-form video at scale for podcasters, social media influencers and digital content creators across multiple platforms with minimal technical expertise. ^{[1] [2] [3]}

Beta testing is scheduled to begin in September 2026 to assess performance, user experience and scalability, with the stated aim of moving toward

commercialization. Michael Mathews, Director of Change Agents, said Phase 2 completion is a key milestone in the company's agentic AI strategy and marks the transition of Catch-Up from a manually configured application into an autonomous AI video studio. The stories do not give a commercial launch date, pricing, beta participant numbers, or any revenue or user figures for the platform. ^{[1] [2] [3]}

Catch-Up sits within a wider portfolio. Change Agents also lists Beacon, an AI search optimization platform, and operates on a Software-as-a-Service model intended to generate recurring subscription revenue. The company distributes KetoAir, a non-invasive consumer breathalyzer measuring ketosis levels, registered with the U.S. Food and Drug Administration as a Class I medical device; story 18 specifies distribution in North America. Separately, Change Agents says it is expanding into AI-enabled autonomous air defense and counter-unmanned aerial systems through its subsidiary Autonomous Air Defense LLC, and is exploring technologies, partnerships and acquisitions in that area. ^{[1] [2] [3]}

For decision-makers, the substance here is a development milestone rather than a product in market. Beta testing is a year out from the Phase 2 announcement by the dates given, so buyers evaluating agentic video tooling have no performance data to work from yet. The breadth of the portfolio, content creation software, search optimization, a consumer medical device and defense systems, is stated in the releases without any disclosed allocation of capital or headcount across those lines. The stories provide no financial results, no customer names and no competitive comparison. ^{[1] [2] [3]}

Who should take heed. Content and marketing leaders assessing agentic video production tools, and investors tracking CHGA, should note this is a pre-beta development milestone with testing not scheduled until September 2026 and no disclosed performance, pricing or customer data.

TECH CORNER

The Agent Data Path: Three Layers, One Missing Primitive

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

The useful framing this week comes from Palo Alto Networks, which splits AI data risk into three distinct layers rather than treating it as one exfiltration

problem. Layer one is the familiar human-to-AI interaction. Layer two is agent-to-tool retrieval over Model Context Protocol, where an agent pulls enterprise data without anyone uploading anything. Layer three is agent-to-agent exchange over A2A, where data moves between agents and never touches a monitored endpoint at all. The argument against conventional DLP is mechanical, not rhetorical: payload inspection matches on content, but the same customer record is legitimate in one flow and a breach in another. Palo Alto's proposed replacement, which it calls Authority-Aware DLP, evaluates user identity, agent identity, intent, sensitivity and destination at the moment of the call and decides inline. The company's implementation is protocol-aware inspection distributed across workforce endpoints, AI runtimes and infrastructure via Prisma Access, Prisma AIRS and its firewalls. The underlying claim is that Zero Trust verification has to be extended from human principals to agent principals, with continuous re-verification rather than a credential issued once. ^[68]

Solo.io is attacking the same problem from the endpoint with agentdesktop, released open source under Apache 2.0. It does four things from a single deployment: real-time discovery of which agents are running and for whom, centralised policy with change previews and continuous reconciliation, identity management that treats each agent as a least-privilege non-human identity with short-lived credentials, and selective telemetry. It integrates with existing identity providers, PKI and LLM gateways, and sits alongside Solo.io's other components: kagent as a Kubernetes-native runtime, agentgateway as the data plane for agent traffic, agentregistry for distribution and agentevals for behavioural evaluation. Note the shape of the design. Short-lived credentials

and continuous reconciliation are the same answer Palo Alto arrives at from the network side, because the failure mode both are addressing is persistent standing access. Rubrik Zero Labs describes that failure mode directly, calling the proliferation of non-human identities tied to agents a shadow workforce, growing faster than governance can track it and carrying persistent access with limited oversight. Two separate 86% figures circulated this week and they are not the same statistic: Solo.io cites 86% of enterprises lacking enforced AI identity access policies, while Rubrik's survey of over 1,600 IT and security leaders found 86% expect agents to outpace their security controls within a year. ^{[65] [63] [68]}

The primitive missing from all of this is reversibility. Rubrik reports that 88% of respondents cannot roll back agent actions without disrupting systems, and that over 80% say agents currently require more manual oversight than the efficiency they return. Detection and policy enforcement stop an action before it happens; neither undoes one that already did. The manufacturing coverage points at the same gap from a different direction, noting that most industrial IT lacks the two-way communication infrastructure needed for automated command and audit, which is what caps how far automation can safely extend regardless of governance intent. Only 23% of Rubrik's respondents claim full visibility over their agents, and Rubrik's own view is that this figure is likely overstated. The practical reading for anyone building here: identity and inline policy are necessary and are now shipping in usable form, but an agent architecture without a rollback path is one where every enforcement decision has to be correct the first time. ^{[63] [64]}

SPONSORSHIP OPPORTUNITY

YOUR ADVERTISEMENT HERE

1600 x 900 OR LARGER · PNG OR JPEG

davidsoden.com/#contact

This space is available.

One slot per issue, your category exclusively, in front of the executives who read this every week.

Enquiries: davidsoden.com/#contact

Briefs

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

Toward Associational Data Rights: A Digital Economy for Families & the Common Good · Government & Public Sector ^[69]

Eliminating wait times - Airport World · Travel, Transportation & Tourism ^[70]

Organizations build software in-house using agentic coding tools instead of buying - ET CIO · Technology & Consumer Electronics ^[71]

India's insurance industry: Can agentic AI build trust? · Banking, Financial Services & Insurance ^[72]

Software Engineering Sr. Manager (GenAI/ Agentic AI/ Fraud & Claims Tech Domain · Banking, Financial Services & Insurance ^[73]

Kyndryl and Google Cloud advance agentic AI at Incore Bank · Banking, Financial Services & Insurance ^[74]

Stop Using AI. Start Commanding It - War on the Rocks · Government & Public Sector ^[75]

Why the next AI hurdle for websites is supporting agents - Diginomica · Technology & Consumer Electronics ^[76]

Huawei Launches Agentic AI Cloud in Nigeria to Accelerate AI Adoption - TechAfrica News · Technology & Consumer Electronics ^[77]

Why Lloyds is Investing PS100m in AI, Skills and Careers | FinTech Magazine · Banking, Financial Services & Insurance ^[78]

OpenClaw 2.0: How to try the updated AI agent, what's new | Mashable · Technology & Consumer Electronics ^[79]

UAH Lab for Applied Science Researchers Engage in International AI Workshop · Government & Public Sector ^[80]

Why Agentic AI is The Next Enterprise Growth Strategy in Australia, According to Bytes Technolab · Banking, Financial Services & Insurance ^[81]

Why Agentic AI is The Next Enterprise Growth Strategy in Australia, According to Bytes Technolab · Banking, Financial Services & Insurance ^[82]

Digital Science's 2026 Catalyst Grant Seeks "Agentic AI Workflows You Can Trust" · Technology & Consumer Electronics ^[83]

Benchmarks Don't Fail Silently. Agentic AI Does. - The Clinical Trial Vanguard · Healthcare ^[84]

From human latency to agentic precision in semiconductor supply chains - Genpact · Technology & Consumer Electronics ^[85]

AI Agents Are Learning to Skirt the Rules. Can Businesses Keep Them Under Control? · Technology & Consumer Electronics ^[86]

Standish Management Expands Technology Ecosystem with DataSnipper's Agentic Platform · Banking, Financial Services & Insurance ^[87]

Build AI skills with new AI-powered learning experiences | oracleuniversity - Oracle Blogs · Technology & Consumer Electronics ^[88]

Microsoft revamps responsible AI rules amid agentic risks - SecurityBrief Australia · Technology & Consumer Electronics ^[89]

The Agentic SOC - From AI Theater to Real Defense - Recorded Future · Technology & Consumer Electronics ^[90]

Netcracker Brings Leadership in Real-World Agentic AI Business Cases for Telecom ... · Media & Communications ^[92]

Turbocharge your canvas development with the canvas authoring agent plugin - Microsoft · Technology & Consumer Electronics ^[91]

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] [Change Agents Completes Phase 2 of its Catch-Up Agentic AI](#) - [GlobeNewswire](#) · [globenewswire.com](#)
- [2] [Change Agents Completes Catch-Up Platform Phase 2 | CHGA Stock News](#) · [stocktitan.net](#)
- [3] [Change Agents Completes Phase 2 of its Catch-Up Agentic AI Video Studio Platform](#) · [aithority.com](#)
- [4] [Fobi AI Launches AgenticBrain, Bringing Conversational Agentic AI to the Fobi Flywheel](#) · [ca.finance.yahoo.com](#)
- [5] [Fobi AI Launches AgenticBrain, Bringing Conversational Agentic AI to the Fobi Flywheel](#) · [finance.yahoo.com](#)
- [6] [Fobi AI Launches AgenticBrain, Bringing Conversational Agentic AI to the Fobi Flywheel](#) · [thestar.com](#)
- [7] [Fobi AI Launches AgenticBrain, Bringing Conversational Agentic AI to the Fobi Flywheel](#) · [aithority.com](#)
- [8] [Genpact Launches Agentic Record-to-Report Suite to Improve Finance Productivity ...](#) · [prnewswire.com](#)
- [9] [Time to get real about outcome-based pricing - can Genpact change the AI risk-reward ...](#) · [diginomica.com](#)
- [10] [AI finance agents: The fastest path to EBITDA recovery - Genpact](#) · [genpact.com](#)
- [11] [Google AI Introduces EnvHarness: A Programmable Layer That Turns Static Agent ... - MarkTechPost](#) · [marktechpost.com](#)
- [12] [Google Mantis: An Agentic Vulnerability Scanning Harness for Reducing False Positives](#) · [infoq.com](#)
- [13] [Google DeepMind Says Gemini Is Evolving From Chatbot To AI Agent](#) · [searchenginejournal.com](#)
- [14] [Tech Mahindra Launches AWS Agentic Process Transformation CoE to Redefine AI-Led ...](#) · [prnewswire.com](#)
- [15] [Hoxton Wealth incorporates agentic AI into wealth management through AWS partnership](#) · [international-adviser.com](#)
- [16] [Clearlake expands AI-fueled tech adoption with Google Cloud - Alternatives Watch](#) · [alternativeswatch.com](#)
- [17] [Globant Introduces MuleSoft AI Pod to Break through Integration Barriers and Scale ...](#) · [prnewswire.com](#)
- [18] [Palo Alto Networks beats quarterly estimates on AI demand, continues acquisition spree](#) · [cnbc.com](#)
- [19] [Palo Alto Networks paid \\$500M for Thrive-backed Console, sources say - TechCrunch](#) · [techcrunch.com](#)
- [20] [CrowdStrike and OpenAI Expand Partnership to Secure the Agentic Era](#) · [ir.crowdstrike.com](#)
- [21] [CrowdStrike and OpenAI Expand Partnership to Secure the Agentic Era](#) · [crowdstrike.com](#)
- [22] [CrowdStrike Fal.Con 2026: CrowdStrike Announces Falcon Guardian and Safe Mind for AIDR](#) · [biztechmagazine.com](#)
- [23] [Introducing Claude Fable 5.1 and Claude Mythos 5.1 - Anthropic](#) · [anthropic.com](#)
- [24] [Who controls the machine? Anthropic extends the MCP vision into the physical world | IBM](#) · [ibm.com](#)
- [25] [Anthropic lays groundwork for bots that shop for you - The Register](#) · [theregister.com](#)
- [26] [Google's Bidding Change Proves Agentic Advertising Has An Agency Problem](#) · [adexchanger.com](#)
- [27] [Agentic AI Ad Platforms in India: Does the ROI Hold Up? - Agency Reporter](#) · [agencyreporter.com](#)
- [28] [Towards Agentic Economies: Eight Cases of New Commerce under AI - Sciences Po](#) · [sciencespo.fr](#)
- [29] [Blackstone leads \\$27 million investment in Israeli cyber startup Huskeys | Ctech](#) · [calcalistech.com](#)
- [30] [Acrab raises US\\$130M amid push to commercialise agentic AI platform | Frontier Enterprise](#) · [frontier-enterprise.com](#)
- [31] [Wonderful Hits \\$5B Valuation as Agentic AI Shifts Toward Embedded Engineering](#) · [forkast.news](#)
- [32] [Halo-record: Open-source audit trails for AI agents - Help Net Security](#) · [helpnetsecurity.com](#)
- [33] [The Database Doesn't Get A Shadow Mode: A DBA's Guide To Agentic AI In Production](#) · [forbes.com](#)
- [34] [The hardest problem in agentic AI begins after the model decides - India Today](#) · [indiatoday.in](#)
- [35] [5 Real-World Applications of Agentic AI in Enterprise Automation - KDnuggets](#) · [kdnuggets.com](#)
- [36] [The Agentic Traceability and Testing Handbook: A Guide to Agentic AI Requirements ...](#) · [militaryembedded.com](#)
- [37] [Beeline Pioneers Agentic AI for the Extended Workforce with Beeline MCP - PR Newswire](#) · [prnewswire.com](#)

- [38] [GK Software Brings Agentic AI to Enterprise Retail - MarTech Cube](#) · [martechcube.com](#)
- [39] [GovCore Announces General Availability of GovCore MCP, Bringing Agentic AI to ... - EIN News](#) · [einnews.com](#)
- [40] [DocuSign Agreement Layer for the Agentic Enterprise Coming to Every Agent - PR Newswire](#) · [prnewswire.com](#)
- [41] [Quiq launches Voice Assist: agentic AI that acts during live calls - PR Newswire](#) · [prnewswire.com](#)
- [42] [Genesys broadens ecosystem for trusted Enterprise Agentic Experience Orchestration](#) · [enterprisetimes.co.uk](#)
- [43] [The Rise of Agentic CCaaS: A New Playbook for Customer Service - Salesforce](#) · [salesforce.com](#)
- [44] [Agentforce G2 Awards: Ranked #1 Best Agentic AI Product | Salesforce CA](#) · [salesforce.com](#)
- [45] [Beyond Basic Chatbots: The Top 7 Agentic AI Platforms Reshaping Customer Support in 2026](#) · [technology.org](#)
- [46] [AI Agents Will Expose the Weakest Links in Your Content Infrastructure - CMS Wire](#) · [cmswire.com](#)
- [47] [AI agents carried out every step of this ransomware attack - then left the victim an 80-page ...](#) · [theregister.com](#)
- [48] [Agentic Ransomware Took Down Enterprise in Ten Hours: AI Left 80-Page Audit](#) · [techtimes.com](#)
- [49] [The First Agentic Attack: How AI Is Reshaping the Economics of Cybersecurity](#) · [forkast.news](#)
- [50] [HPE raises FY27 outlook on agentic AI demand, reports record orders - Yahoo Finance](#) · [finance.yahoo.com](#)
- [51] [Arm Holdings Rebounds as Agentic AI Demand Puts \\$258 Breakout in Focus - TradingKey](#) · [tradingkey.com](#)
- [52] [Hot Chips 2026: SN50 RDU Dataflow at Scale - SambaNova](#) · [sambanova.ai](#)
- [53] [NVIDIA and MediaTek Deepen Long-Standing Partnership to Build AI Edge to Cloud ...](#) · [nvidianews.nvidia.com](#)
- [54] [Broadcom Delivers End-to-End Security, Identity, and Observability for Agentic AI](#) · [stocktitan.net](#)
- [55] [Broadcom Unveils Security, Identity and Observability Capabilities for Agentic AI on Private AI Cloud](#) · [quiverquant.com](#)
- [56] [Broadcom Delivers End-to-End Security, Identity, and Observability for Agentic AI](#) · [markets.businessinsider.com](#)
- [57] [Private AI cloud, agentic infrastructure dominate VMware Explore](#) · [networkworld.com](#)
- [58] [Broadcom Unveils AgentMinder, An Enterprise Solution for AI Agent Governance and ...](#) · [finance.yahoo.com](#)
- [59] [Scaling the Agentic Enterprise: Production-ready AI Agents with Tanzu Platform](#) · [blogs.vmware.com](#)
- [60] [What "agentic" actually means for software quality \(beyond the buzzwords\) - DevPro Journal](#) · [devprojournal.com](#)
- [61] [GeekyAnts Joins AgentsNexus India 2026 as Agentic AI Moves Beyond the Demo Stage](#) · [wboc.com](#)
- [62] [UAE Enterprises Embrace Agentic AI as Governance and Architecture Become Critical - Pune Mirror](#) · [punemirror.com](#)
- [63] [As agentic AI adoption accelerates, Rubrik warns of growing security gaps - Tahawul Tech](#) · [tahawultech.com](#)
- [64] [How Most Manufacturers Are Adopting Agentic AI Without A Governance Plan](#) · [roboticstomorrow.com](#)
- [65] [Solo.io Extends Agentic Governance to the Desktop with agentdesktop - AiThORITY](#) · [aithority.com](#)
- [66] [How poor CRM data continues to plague Marketing's AI initiatives, but the board may not see it](#) · [diginomica.com](#)
- [67] [Agentic AI won't fix bad subscription data - SDxCentral](#) · [sdxcentral.com](#)
- [68] [Securing the Agentic Shift: Data Protection Across the AI Data Path - Palo Alto Networks](#) · [paloaltonetworks.com](#)
- [69] [Toward Associational Data Rights: A Digital Economy for Families & the Common Good](#) · [ifstudies.org](#)
- [70] [Eliminating wait times - Airport World](#) · [airport-world.com](#)
- [71] [Organizations build software in-house using agentic coding tools instead of buying - ET CIO](#) · [cio.economictimes.indiatimes.com](#)
- [72] [India's insurance industry: Can agentic AI build trust?](#) · [insuranceasia.com](#)
- [73] [Software Engineering Sr. Manager \(GenAI/ Agentic AI/ Fraud & Claims Tech Domain](#) · [wellsfargojobs.com](#)
- [74] [Kyndryl and Google Cloud advance agentic AI at Incore Bank](#) · [kyndryl.com](#)
- [75] [Stop Using AI. Start Commanding It - War on the Rocks](#) · [warontherocks.com](#)
- [76] [Why the next AI hurdle for websites is supporting agents - Diginomica](#) · [diginomica.com](#)
- [77] [Huawei Launches Agentic AI Cloud in Nigeria to Accelerate AI Adoption - TechAfrica News](#) · [techafricanews.com](#)
- [78] [Why Lloyds is Investing PS100m in AI, Skills and Careers | FinTech Magazine](#) · [fintechmagazine.com](#)
- [79] [OpenClaw 2.0: How to try the updated AI agent, what's new | Mashable](#) · [mashable.com](#)
- [80] [UAH Lab for Applied Science Researchers Engage in International AI Workshop](#) · [uah.edu](#)
- [81] [Why Agentic AI is The Next Enterprise Growth Strategy in Australia, According to Bytes Technolab](#) · [finance.yahoo.com](#)
- [82] [Why Agentic AI is The Next Enterprise Growth Strategy in Australia, According to Bytes Technolab](#) · [wboc.com](#)

- [83] Digital Science's 2026 Catalyst Grant Seeks "Agentic AI Workflows You Can Trust" · finance.yahoo.com
- [84] Benchmarks Don't Fail Silently. Agentic AI Does. - The Clinical Trial Vanguard · clinicaltrialvanguard.com
- [85] From human latency to agentic precision in semiconductor supply chains - Genpact · genpact.com
- [86] AI Agents Are Learning to Skirt the Rules. Can Businesses Keep Them Under Control? · news.darden.virginia.edu
- [87] Standish Management Expands Technology Ecosystem with DataSnipper's Agentic Platform · prnewswire.com
- [88] Build AI skills with new AI-powered learning experiences | oracleuniversity - Oracle Blogs · blogs.oracle.com
- [89] Microsoft revamps responsible AI rules amid agentic risks - SecurityBrief Australia · securitybrief.com.au
- [90] The Agentic SOC - From AI Theater to Real Defense - Recorded Future · recordedfuture.com
- [91] Turbocharge your canvas development with the canvas authoring agent plugin - Microsoft · microsoft.com
- [92] Netcracker Brings Leadership in Real-World Agentic AI Business Cases for Telecom ... · finance.yahoo.com