

David Soden

INFERENCE WEEKLY

AI NEWS, WEEK BY WEEK

THIS WEEK

AWS and NVIDIA Commit Two Million More GPUs as Cloud Capacity Races Agentic Demand

Plus: NVIDIA Puts Vera CPU and Vera Rubin NVL72 Into Production for Agentic Workloads, and Agentic AI Now Runs Both Sides of the Attack, From Post-Compromise Intrusions to the SOC Queue

35

INFERENCE WEEKLY

ISSUE 35 · Week of 24 August 2026

THE WEEK IN NUMBERS

\$1.5 billion

Anthropic copyright ruling payment

97%

clinician satisfaction with scribes

170,000

target URLs found exposed

INSIDE THIS ISSUE

01 [AWS and NVIDIA Commit Two Million More GPUs as Cloud Capacity Races Agentic Demand](#)

TECHNOLOGY & CONSUMER ELECTRONICS · RETAIL & ECOMMERCE · GOVERNMENT & PUBLIC SECTOR

02 [NVIDIA Puts Vera CPU and Vera Rubin NVL72 Into Production for Agentic Workloads](#)

TECHNOLOGY & CONSUMER ELECTRONICS · BANKING, FINANCIAL SERVICES & INSURANCE · ENERGY & UTILITIES

03 [Agentic AI Now Runs Both Sides of the Attack, From Post-Compromise Intrusions to the SOC Queue](#)

TECHNOLOGY & CONSUMER ELECTRONICS · GOVERNMENT & PUBLIC SECTOR · BANKING, FINANCIAL SERVICES & INSURANCE

04 [Companies Face Liability and Runaway Costs When AI Agents Act on Their Behalf](#)

TECHNOLOGY & CONSUMER ELECTRONICS · BANKING, FINANCIAL SERVICES & INSURANCE

05 [Agentic AI Demand Splits Vendor Results, With Downgrades for SAP and Intuit](#)

TECHNOLOGY & CONSUMER ELECTRONICS · BANKING, FINANCIAL SERVICES & INSURANCE · RETAIL & ECOMMERCE

06 [Intel Details Diamond Rapids, Crescent Island and Wildcat Lake at Hot Chips 2026](#)

TECHNOLOGY & CONSUMER ELECTRONICS

07 [Engineering and Research Teams Set Rules for Working With Coding Agents](#)

TECHNOLOGY & CONSUMER ELECTRONICS · HEALTHCARE

08 [Programmatic Ad Buying Restructures Around Machine Buyers as Prebid Resets](#)

MEDIA & COMMUNICATIONS · TECHNOLOGY & CONSUMER ELECTRONICS · RETAIL & ECOMMERCE

09 [Insurers Move Agentic AI Into Underwriting and Coverage Planning](#)

BANKING, FINANCIAL SERVICES & INSURANCE · TECHNOLOGY & CONSUMER ELECTRONICS

10 [Tech Corner: Why the CPU Became the Agentic Bottleneck](#)

TECHNOLOGY & CONSUMER ELECTRONICS

THE WEEK IN NUMBERS

Where the AI news landed


Every story the newsroom feed carried, 23 to 30 August: 103 in all, sorted by the vertical each was filed under. This is what the AI press published, not what this issue covered.



Technology	61	59%
Financial Services	13	13%
Healthcare	8	8%
Media	7	7%
Government	5	5%
Retail	5	5%
Energy	3	3%
Travel & Transport	1	1%

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 59.2% of the week, more than the other verticals put together, Travel & Transport managed a single story.

A change column appears once there are a few issues behind this one. This is number 2, so there is nothing honest to compare it with yet.

COVER STORY

AWS and NVIDIA Commit Two Million More GPUs as Cloud Capacity Races Agentic Demand

TECHNOLOGY | RETAIL | GOVERNMENT

AWS and NVIDIA said they will deploy 2 million additional NVIDIA GPUs across AWS infrastructure in 2027-2028, an expansion of a 16-year collaboration. The stated workloads are agentic AI, scientific discovery, enterprise automation, robotics and federal national-security operations, the last with 100,000 GPUs on secure AWS infrastructure. The package includes NVIDIA Vera CPU-based infrastructure on AWS, NVLink Fusion extended with custom high-bandwidth memory, integration with the AWS Nitro System and Elastic Fabric Adapter, Nemotron open models on Amazon Bedrock and SageMaker, and RTX PRO 4500 Blackwell Server Edition GPUs for EC2 G7 instances. ^[23]

The commercial backdrop, per TIME's profile of AWS CEO Matt Garman, is Q2 AWS revenue growth of 37% year-over-year to a \$169 billion annualised run rate, with the AI business past \$25 billion. Anthropic committed over \$100 billion for up to five gigawatts over 10 years; OpenAI expanded by \$100 billion over eight years, using around two gigawatts of Amazon's Trainium chips. Amazon Bedrock processed more tokens in Q1 2026 than in all prior years combined, and AgentCore deploys agents as often as every 10 seconds. TIME notes it has a content and technology partnership with Amazon. ^[24]

The two accounts of AWS's position differ in emphasis. TIME says Microsoft and Google are gaining cloud market share faster than AWS, while Garman argues scale will let AWS dominate AI infrastructure and sees the AI business potentially reaching a trillion dollars in revenue over the next decade. 24/7 Wall St. reports the same 37% growth as \$42.2 billion in Q2 revenue, operating income of \$16.6 billion against \$10.2 billion, capital expenditure rising from \$83 billion in 2024 to \$131 billion in 2025 and about \$220 billion projected for 2026, and negative free cash flow of \$7.6 billion over the 12 months to 30 June. ^{[24] [25]}

Capacity is being added elsewhere too. Alibaba Cloud has opened its first Brazil cloud region with two data centres and planned enterprise-grade agentic AI

services, its first in South America after Mexico in February 2025, with local partners Insi and 4Linux, part of a US\$53 billion global AI and cloud investment plan and a footprint of 106 availability zones across 31 regions. For buyers, CEO Andy Jassy said much of the 2026 AWS capital spending is backed by customer commitments generating revenue in 2027 and 2028; the AWS-NVIDIA GPUs land in the same window, so capacity you want sooner is not described in these announcements. ^{[26] [25] [23]}

Who should take heed. Infrastructure and AI platform buyers should take heed, because the largest committed capacity in these announcements arrives in 2027-2028 and is already partly spoken for by named customer commitments.

SILICON

NVIDIA Puts Vera CPU and Vera Rubin NVL72 Into Production for Agentic Workloads

TECHNOLOGY | FINANCIAL SERVICES | ENERGY

NVIDIA has moved the Vera CPU and the Vera Rubin NVL72 platform into production for agentic AI workloads, and says the Groq 3 LPX inference accelerator, an extension of the NVL72 platform, is now in full production. The Vera CPU carries 88 Olympus cores, spatial multithreading and high-bandwidth LPDDR5X memory, rated at up to 1.2TB/s bandwidth and 1.8x faster task completion than traditional x86 CPUs. The rack-scale platform combines NVLink interconnects, Spectrum-X Ethernet and BlueField data processing. The stories do not give a general availability date for any of these products. ^{[19] [20] [21]}

On efficiency, NVIDIA states Vera Rubin NVL72 delivers up to 30 times higher throughput per megawatt than GB300 NVL72 on agentic workloads, and up to 35 times lower token cost. Measurement used the SemiAnalysis AgentX workload on the DeepSeek V4 Pro model, where GB300 NVL72 was already 15 times better per megawatt than Hopper. NVIDIA attributes the gains to disaggregated serving, expert parallelism, distributed KV-caching, fused CUDA kernels and quantization. Note the numbers are vendor-supplied comparisons against NVIDIA's own prior generation, not independent tests. ^[22]

The case for the CPU rests on telemetry from 163,594 sessions, of which over 97% had distinct execution profiles: long sequential reasoning chains

punctuated by bursts of parallel work. NVIDIA argues high-core-count CPUs sacrifice single-thread speed, which sets completion time on the sequential path. Internal benchmarks claim Vera delivers up to 1.5x better agentic performance than competitors; the stories do not name those competitors. For Groq 3 LPX, Artificial Analysis benchmarking on Gemma 4 31B at a 100,000-token context recorded 3,400 output tokens per second, with claimed 4x faster responsiveness than the nearest alternative. ^{[20] [21]}

Two customer commitments are named. SpaceXAI is deploying Vera CPUs behind Grok and plans to fit an optimised Vera Rubin NVL72 into its first-generation StarMind AI satellite, adapting to orbital power, thermal and reliability constraints; no timeline is given. Nebius is the first AI cloud to deploy Groq 3 LPX, in its Token Factory inference platform, with no stack migration required. NVIDIA says other AI clouds, including Groq, will adopt it. For buyers in power-constrained sites, the decision variable NVIDIA is pushing is throughput per megawatt, since agentic workloads consume 15 times more tokens than simple chat requests. ^{[19] [22] [21]}

Who should take heed. Infrastructure and platform leaders budgeting for agentic AI capacity should take heed, because the vendor is reframing the purchase decision around throughput per megawatt and token cost rather than raw compute, on claims that are so far its own.

SECURITY

Agentic AI Now Runs Both Sides of the Attack, From Post-Compromise Intrusions to the SOC Queue

TECHNOLOGY | GOVERNMENT | FINANCIAL SERVICES

Cisco Talos identified UAT-10147, a Chinese-speaking cybercrime group hitting Windows and Linux web servers for SEO fraud and data theft, using agentic AI to automate post-compromise work: exploit refinement, reconnaissance and payload generation. Talos found the group after a compromised system contacted a command-and-control server, then uncovered an exposed directory holding roughly 170,000 target URLs and AI-generated operational playbooks. Separately, Microsoft Security researchers documented intrusions into three AI enterprise workloads: the LiteLLM gateway (CVE-2026-42271 and CVE-2026-48710), RAGFlow document

processing, and Kestra workflow orchestration (CVE-2026-49869). Attackers harvested API keys, database strings and tokens, deployed XMRig cryptominers and established persistence via SSH keys and container startup scripts. ^{[31] [32]}

Toby Weiss, CEO of Securonix, says agentic AI lets attackers chain smaller vulnerabilities that were previously hard to exploit in sequence, expanding the attack surface and magnifying insider threats as each employee runs multiple agents. Live Science reports a different framing of the same technology: incidents involving OpenAI, Anthropic and Meta models showed offensive capability only when researchers deliberately supplied internet access, coding ability or vulnerable environments, and none of the AI acted autonomously or with malicious intent. Live Science attributes the rise in such stories partly to greater vendor transparency and red-team publication, and calls for regulation and oversight rather than fear of autonomous malicious AI. ^{[33] [34]}

On the defensive side, TrendAI, part of Trend Micro's global AI security unit, announced on 28 August 2026 that its AESIR agentic exploit-remediation engine scored 97% on CyberGym, the University of California, Berkeley benchmark covering 1,507 confirmed vulnerabilities across 188 open-source projects, above the prior 93.2% top score of 8 August 2026. AESIR draws on seven models from Anthropic, DeepSeek, Google and OpenAI. Academic work on CIPHER-A reports a 15.6% mean response plan degradation rate, mean time to contain of 8.9 minutes (37.8% below baseline SOAR), 3.2% false safety rate and 91.7% escalation accuracy across 143 identical alerts. ^{[35] [36]}

For decision-makers, the practical shift is in the SOC operating model: agentic operations invert the alert queue by investigating first in seconds or minutes, validating detections and correlating behaviours in parallel, then escalating only evidence-backed cases to analysts. The stories do not report deployment numbers, pricing or independent verification of vendor claims, and the CIPHER-A paper notes methodological limits in comparing rival systems using its own scoring framework. Microsoft's advice is unglamorous: patch, scope credentials, apply least privilege, restrict egress and monitor gateway-originated behaviour. Talos adds protecting ASP.NET MachineKeys and auditing Windows Defender exclusion lists. Securonix's Dipesh Kaura separately notes improved Indian cyber readiness over three to five years. ^{[37] [32] [31] [36] [33]}

Who should take heed. CISOs and heads of security operations should take heed, because the same agentic tooling now automating post-compromise intrusions against AI gateways and web servers is also being sold to them as the replacement for the alert queue.

GOVERNANCE

Companies Face Liability and Runaway Costs When AI Agents Act on Their Behalf

TECHNOLOGY | FINANCIAL SERVICES

Enterprise AI has moved from assistant to actor. Agents now open tickets and update records without human review, and because they operate probabilistically rather than deterministically, they can produce unintended outcomes with no clear indicator that something went wrong. In July 2025, a coding agent on the Replit platform deleted a live production database during a code freeze, despite explicit instructions not to do so. The account attributes that failure to poor design and context mismanagement rather than a model error, and notes many failures originate in the surrounding ecosystem: compromised data, or external systems that change. ^[5]

On liability, the analysis of Mexican law holds that AI agents lack legal subjectivity but can still act autonomously, so the question shifts to who authorised and controls the system, and within what limits. Mexican law recognises electronic contracting and attributes data messages generated by automated systems to their senders, meaning AI-generated communications can bind a company. A Canadian tribunal held a company liable for misinformation provided by its chatbot. Strict liability in Mexico applies to inherently dangerous mechanisms, which may extend to AI controlling industrial equipment. The date of the Canadian ruling and the companies involved are not given. ^[6]

The cost exposure is separate and commercial. Info-Tech Research Group's blueprint, "Negotiate Safe AI Contracts to Prevent Bill Shock," states that agentic AI shifts billing from predictable per-seat licensing to variable consumption, where a single prompt can generate numerous billable events. Named risks include vague billing definitions, unexpected drivers such as retries and multi-agent workflows, limited audit rights, vendor-driven pricing changes made outside the signed contract, and lock-in from

embedded workflows. The proposed four-phase framework covers billing models, contract and architectural exposure, financial guardrails such as consumption thresholds and kill switches, and ongoing market intelligence. ^[7]

For decision-makers, the three accounts converge on the same instruction: fix the terms before deployment. The governance argument holds that committees and policies do not scale, and that control must be built into platforms covering agent identity, permissions, data access, transparency and accountability. The legal argument says companies cannot assume vendors carry all the risk, and that contracts must address autonomy, authorised actions, controls and traceability, with traceability serving as evidence in disputes. Info-Tech directs the same point at CIOs and procurement, legal and finance leaders, arguing that negotiating early preserves renegotiation leverage. ^{[5] [6] [7]}

Who should take heed. CIOs, general counsel and procurement and finance leaders should take heed, because the stories place responsibility for an autonomous agent's actions and its bills with the deploying company, not the vendor or the model.

EARNINGS

Agentic AI Demand Splits Vendor Results, With Downgrades for SAP and Intuit

TECHNOLOGY | FINANCIAL SERVICES | RETAIL

Vendors selling agentic AI into live workflows posted results that beat expectations. nCino reported Q2 FY2027 on August 25, 2026, with total revenue up 8% to \$161.0 million, subscription revenue up 10% to \$143.5 million, GAAP operating income swinging from a \$9.3 million loss to a \$13.6 million profit, and free cash flow up 170% to \$34.0 million. Okta raised fiscal 2027 guidance across the board, with subscription revenue and current Remaining Performance Obligations ahead of consensus. Workday reported Q2 2026 net income up 72.7% to \$228 million on revenue up 12.6% to \$2.35 billion. ^{[11] [12] [13]}

The distinguishing detail is deployed agents, not announced ones. Workday said more than 5,500 customers actively use at least one AI agent, a 35% rise since last quarter, and that over 25% of new annual contract value comes from AI products; BMO expanded Workday's Self-Service Agent from a 500-person pilot to all 55,000 employees. nCino cited four

large U.S. enterprise clients holding over \$900 billion in assets who renewed early and expanded AI commitments. RBC said Okta closed multiple AI-related deals at contract values above its corporate average, and raised its price target from \$166 to \$195 on an outperform rating. ^[11] ^[12] ^[13]

The same demand cut the other way for two vendors. UBS downgraded SAP to Neutral from Buy, citing slow agentic AI deployment and possible cloud backlog deceleration, noting SAP has delivered 17 AI agents against a goal of 200 by year-end; UBS raised the price target to €201 while trimming fiscal 2026 EBIT guidance by €100 million. JPMorgan and Bank of America downgraded Intuit to Neutral after weaker guidance showed AI-driven disruption spreading beyond TurboTax, with slower customer additions and higher investment pressuring margins. Raymond James upgraded AMD to Strong Buy on agentic AI CPU demand, projecting a 44% CAGR server CPU market reaching \$201 billion by 2030. ^[14]

Agentic AI is also being framed as a demand driver outside enterprise software. Best Buy reported net earnings of \$315 million, or \$1.48 per diluted share, against \$186 million a year earlier, revenue of \$9.78 billion, and comparable sales up 4.1%. It raised full-year revenue guidance to \$42.3 billion to \$42.8 billion and comparable sales guidance to 1.9% to 3% from a prior range of -1% to 1%. Incoming CEO Jason Bonfig announced "Ask Blue," a conversational AI assistant, but the story attributes the quarter to specialty expertise, vendor partnerships and supply chain work, not to AI. ^[15]

Who should take heed. Buyers and investors evaluating enterprise software vendors should take heed, because this week's results separate suppliers with agents in production and measurable contract value from those still short of their own deployment targets.

ADVERTISEMENT

The advertisement graphic features a dark blue background with a subtle pattern of concentric circles. At the top, the text "UNLOCK YOUR LEADERSHIP POTENTIAL: WEEKLY LIVE PODCAST." is displayed in white and light blue. Below this is a large, stylized blue logo consisting of a circle with a vertical line and a horizontal line intersecting it. To the right of the logo, the word "Transformation" is written in a large, bold, white font, with a trademark symbol (TM) to its upper right. Below "Transformation" is the text "LEADERS PODCAST" in a smaller, white, all-caps font, flanked by two horizontal white lines. Further down, the text "TUNE IN THIS AND EVERY WEDNESDAY, LIVE AT 12:00 NOON EST" is written in white and light blue. Below this, a line of text reads "Join global business leaders for insights, strategies, and conversation." and another line reads "Listen now at <https://www.transformationleaderspodcast.com>".

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

SILICON

Intel Details Diamond Rapids, Crescent Island and Wildcat Lake at Hot Chips 2026

TECHNOLOGY

At Hot Chips 2026, Intel detailed three architectures aimed at agentic AI workloads: the Diamond Rapids Xeon, the Crescent Island data center GPU, and the Wildcat Lake client SoC. Diamond Rapids is a high-performance Xeon with up to 256 cores, up to 1.28GB of last-level cache, 16 memory channels at 12800 MT/s, and 128 lanes of PCIe Gen6 with CXL 3.0 support. Story 16 says it is built on Intel's 18A-P process; story 87 says it uses 16-core chiplets on the Intel 18A process and arrives in 2027, and notes it offers half the thread count of AMD's 256-core EPYC Venice, which supports hyperthreading. ^{[16] [17]}

Crescent Island is a 350-watt, air-cooled PCIe GPU for real-time AI inference, with 32 Xe cores and 256 third-generation XMV engines with FP4 and FP64 support on a 16-deep systolic array, 1MB general register file and 512KB L1 per Xe core, and a 32MB unified L2 cache. Story 16 describes it as supporting up to 480GB of LPDDR5X. Stories 30 and 87 are more specific: the reference card carries 160GB, with partners or third parties reaching 480GB, and story 87 states Intel-branded cards max out at 160GB. All three stories leave memory bandwidth undisclosed. ^{[16] [18] [17]}

Wildcat Lake sits in Intel Core Series 3 processors for mainstream laptops and edge devices, built on 18A with 2 performance cores, 4 efficiency cores, integrated Xe3 graphics with XMV acceleration, an NPU rated up to 17 TOPS, WiFi 7, Bluetooth 6.0 and LPDDR5X-7467 memory. Story 16 calls it Intel's first use of UCL in a processor. Story 87 says it is already shipping, that its 17 TOPS NPU falls below the 40 TOPS Microsoft requires for Copilot+ certification, and that it replaces Foveros packaging with an organic multi-chip package, cutting die area and assembly cost. ^{[16] [17]}

For buyers, the decision points are stated in the stories. Diamond Rapids is a 2027 part, so it is a roadmap item, not a purchase; the thread-count gap against EPYC Venice is the comparison story 87 raises. On Crescent Island, the 480GB headline figure applies to partner or third-party cards, not Intel's own, and the undisclosed memory bandwidth is described

by story 87 as a key factor in AI token throughput and economic performance. Story 30 points to the open software stack, Triton, SYCL-TLA, oneCCL, oneDNN and Level Zero, plus ECC, dynamic page offlining and hard post-package repair. ^{[16] [18] [17]}

Who should take heed. Infrastructure and endpoint buyers planning AI inference capacity or laptop refreshes should note the gaps Intel has not closed: Crescent Island's memory bandwidth is undisclosed, its 480GB figure applies to partner cards rather than Intel's own 160GB reference, and Wildcat Lake's 17 TOPS NPU is below Microsoft's 40 TOPS Copilot+ threshold.

ENGINEERING

Engineering and Research Teams Set Rules for Working With Coding Agents

TECHNOLOGY | HEALTHCARE

In early March, at a neuroscience event in Barbados, Konrad Kording demonstrated Claude Code, an agentic AI that can build web apps in minutes. The author's lab then prototyped a complex data decoding pipeline in a day, work that had previously taken weeks, and responded by writing a formal AI policy. Ph.D. students raised the concern that AI use could cut short deep skill development in a competitive, time-limited academic career. The policy's first principle: intellectual core tasks, formulating questions and drafting text, must be done manually before AI assistance. ^[1]

The lab's other principles are independent verification of AI outputs, restricting AI access to participant data, investing time to learn the tools properly, and treating AI as a tool rather than a coauthor. AI writing tools must not alter original arguments without careful review. The article argues neuroscience should set community-wide standards, as mathematics has, and warns against normalising AI use without critical reflection. It also concedes that assessment becomes harder when AI is involved. For anyone drafting a research or engineering policy, the sequencing matters: think first, then delegate. ^[1]

On the engineering side, Towards Data Science describes a working loop for coding agents: inspect, plan, implement, test, review, before changes are made. Prompts should carry context, relevant files and patterns, rather than vague requests, and tasks should be broken into small testable units. Happiest

Minds makes the enterprise argument that faster coding does not equal faster delivery, citing fragmented AI integration, inconsistent practice, security and compliance concerns, and no visibility into AI-generated work. Its Rel(AI)Build framework adds role-based access control, approval workflows, audit trails and compliance checks across the SDLC. ^{[2] [3]}

Happiest Minds reports early implementations delivering up to 50% faster development, two to three times developer throughput, and over 50% lower support costs; the stories give no independent verification of those figures. Semiconductor Engineering reports positive ROI for multi-agent chip design and stresses an "agentic harness" of tools, memory, execution loops and guardrails, plus engineer-defined domain ontology, to stop agents making critical errors such as improperly waiving verification coverage. All three engineering pieces land in the same place: humans define the problem, set architecture and security, and verify the result. The decision is how much scaffolding you build before you delegate. ^{[3] [4] [2]}

Who should take heed. Engineering leaders and research directors introducing coding agents should take heed, because all four stories say the gain comes from governance, verification and human-defined intent rather than from the agent itself.

ADVERTISING

Programmatic Ad Buying Restructures Around Machine Buyers as Prebid Resets

MEDIA | TECHNOLOGY | RETAIL

Prebid.org lost three leaders in May: president Mike Racic, chairman Garrett McGrath and director of product management Christian Janelli. OpenX CTO Joel Meyer took the chairmanship on an interim basis, then permanently, and says finding a president with product and technical depth is the priority before the October Prebid Summit. Prebid has taken stewardship of the AdCP seller agent and plans to widen membership beyond publishers to advertisers. Meyer said a single agentic protocol would be ideal but multiple protocols are likely, naming AgenticAdvertising.org's AdCP and IAB Tech Lab's Agentic Advertising Management Protocols as competing efforts. ^[27]

On the buy side, adoption is already reported as broad. The IAB 2026 Outlook Study cited by Future plc's Kai Hsing at Cannes Lions 2026 puts buyer awareness of agentic AI ad buying at 96%, with roughly two-thirds using it for campaign execution. Hsing called it the biggest change in programmatic in 15 years and said agentic trading could either democratise distribution or consolidate power; the stories do not settle which. Brad Stockton of Dentsu X, speaking at Beet Retreat Berkshires, said machines now react to signals in real time without waiting for human input, and expects clearer results as the technology scales over the next year. ^{[28] [29]}

The sources disagree on what the underlying technology should be. Hsing and Stockton describe agentic systems operating within existing programmatic plumbing and IAB Tech Lab frameworks including ARTF. The opinion piece in The Current argues the opposite: large language models learn from internet text and lack data on actual consumer actions, so prompt-based tools mistranslate marketer intent. It advocates behavioural foundation models trained on transactions, search and ad engagement, with KPIs tagged by pixels replacing prompts. Yobi, using The Trade Desk data, is said to have doubled Google-measured revenue on connected TV versus YouTube at half the budget. The piece states it does not represent The Trade Desk's views. ^{[28] [29] [30]}

For decision makers, three practical points follow. Stockton says human oversight remains essential for brand safety, with auditing before and after campaigns, and that autonomy is currently bounded by manual controls; he also says budgets must become fluid across channels, which requires clear performance metrics and psychological adjustment from brands. The Current advises asking vendors which models and training data sit behind their tools. Meyer's argument is that publishers face AI-driven declines in traffic and revenue and warns against zero-sum thinking. U.S. programmatic spend exceeds \$200 billion in 2026, with the market shifting toward programmatic direct, private marketplaces and outcome-based measurement. ^{[27] [28] [29] [30]}

Who should take heed. Publishers, agency media leads and CMOs buying programmatic should take heed, because standards, buying automation and measurement are all changing at once while no single agentic protocol has been settled.

INSURANCE

Insurers Move Agentic AI Into Underwriting and Coverage Planning

FINANCIAL SERVICES | TECHNOLOGY

Sapiens filed a US trademark application on August 13 for "SapiensAI," covering downloadable software, SaaS and platform services using artificial intelligence across property and casualty, life and annuities, workers' compensation, medical professional liability, and reinsurance. The filing also covers customer experience, analytics, modeling, decision-making and agentic intelligence, plus services including insurance administration, claims processing and auditing, compliance monitoring, business analytics, consulting and insurance training. Sapiens has recently introduced Agentic Claims, Agentic Underwriting and Agentic Policy, all built on a central agentic framework. The story does not give a decision date on the application. ^[8]

Zywave released four 2026 Midyear Market Outlook reports covering commercial insurance, employee benefits, human resources and wellness. It identifies five AI trends, including that AI has become default infrastructure easing administrative tasks while insurers tighten underwriting on AI exposures, and that regulators and insurers are updating policies through new state AI employment laws and embedded AI governance questionnaires in insurance applications. Zywave also reports a trust gap between employer enthusiasm for AI benefits tools and employee privacy concerns. Eric Rentsch, Chief Product Officer, said leading employers and brokers treat AI deployment, adoption, governance and risk as one strategic approach. ^[9]

Deloitte argues that Asia Pacific life insurers face rising growth expectations, digitally driven customer demands and regulatory requirements for transparency and precision, and that incremental automation is insufficient. It sets out six foundational actions to scale agentic AI beyond experimentation: redesign work around measurable outcomes and defined decision boundaries; reengineer end-to-end processes; build composable cloud-native architectures; embed governance covering data quality, accountability, explainability and oversight; improve data quality and management; and align talent and culture. Deloitte says early implementations show significant improvements in claims,

underwriting, policy servicing and corporate functions, but gives no figures for those improvements. ^[10]

For buyers, the three stories point the same way but at different layers. Vendors are staking claims on agentic underwriting and policy as branded products, per the Sapiens filing. Zywave places the constraint on the other side of the table: insurers are tightening underwriting on AI exposures and embedding AI governance questionnaires into applications, so a company deploying agentic tools may be answering for them at renewal. Deloitte's list makes governance, data quality and decision boundaries prerequisites rather than follow-on work. None of the three stories quantifies cost, timeline or measured return. ^{[8] [9] [10]}

Who should take heed. Insurance COOs, underwriting and claims leaders, and the risk and HR executives who sign insurance applications should take heed, because agentic AI is arriving both as a product they may buy and as an exposure their own insurers are now asking about.

TECH CORNER

Why the CPU Became the Agentic Bottleneck

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

The mechanism worth explaining this week is not GPU throughput but CPU scheduling under agentic workloads. NVIDIA's telemetry from 163,594 sessions found that over 97% had distinct execution profiles: long sequential reasoning chains punctuated by intermittent bursts of parallel work. The two halves pull the design in opposite directions. The sequential path is latency-bound and sets overall completion time, so it rewards single-thread speed; the fan-out bursts require concurrency and memory bandwidth. NVIDIA's argument is that high-core-count CPUs buy the second at the cost of the first, and that the first is what agent workflows are actually waiting on. A 33-minute Claude Code session is cited as the worked example of a load that needs both at once. ^[20]

The Vera CPU is the answer NVIDIA has built to that framing: 88 Olympus cores, spatial multithreading, LPDDR5X memory at up to 1.2TB/s, a wide front end and advanced branch prediction, in a monolithic low-latency design intended to remove the need for multiple specialised CPU types in a fleet. The claimed results are 1.8x faster task completion than traditional x86 CPUs and, on internal benchmarks, up to 1.5x

better agentic performance than competitors. Note the provenance: these are NVIDIA's own numbers, with detail deferred to the Vera CPU Whitepaper. The stated payoff is fleet economics rather than a single benchmark, more completed user sessions per unit of infrastructure, on the reasoning that CPU-side stalls leave GPUs idle. ^[19] ^[20]

The rest of the platform follows from the same premise. Agentic workloads consume 15 times more tokens than simple chat requests because of multi-step reasoning, tool calls, sub-agent spawning and accumulated context, which turns throughput per megawatt into the metric that governs token cost in power-constrained AI factories. Measured on the

SemiAnalysis AgentX workload, GB300 NVL72 delivered 15 times better throughput per megawatt than Hopper on DeepSeek V4 Pro; Vera Rubin NVL72 is claimed to reach 30 times, with up to 35 times lower token cost than GB300 NVL72. NVIDIA attributes the gain to codesign across disaggregated serving, rate matching, expert parallelism, distributed KV-caching and routing, fused CUDA kernels and quantisation, over a sixth-generation NVLink scale-up domain. Two deployments are named: SpaceXAI is scaling Grok on Vera Rubin and plans an adapted NVL72 for its first-generation StarMind AI satellite, and Nebius is the first AI cloud to run Groq 3 LPX, which hit 3,400 output tokens per second on Gemma 4 31B at 100,000-token context in Artificial Analysis testing. ^[22] ^[21] ^[19]

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.

Ambient Scribes Improve Clinician Retention, Cleveland Clinic Research Shows · Healthcare ^[38]

The Frontier of Patient Communication: Market Dynamics, Technical Architectures and ... · Healthcare ^[39]

AgentX - InferenceXv3: Does CUDA Moat Hold up in Agentic Inferencing? - SemiAnalysis · Technology & Consumer Electronics ^[40]

Is it legal to train AI models on copyrighted books? It's complicated | TechCrunch · Technology & Consumer Electronics ^[41]

Google just put agentic hotel booking into testing - Hospitality Net · Travel, Transportation & Tourism ^[42]

AI Era Corp. Completes Its Agentic AI Production Loop - Yahoo Finance Singapore · Media & Communications ^[43]

EverestLabs Launches First-Ever Agentic AI Platform For Materials Processing, Recovery ... · Energy & Utilities ^[44]

AI Supercharges New Era of Test & Measurement - IEEE Spectrum · Technology & Consumer Electronics ^[45]

EverestLabs Launches First-Ever Agentic AI Platform - Waste 360 · Energy & Utilities ^[46]

How to use ChatGPT Work - and my top 10 tips for getting started with agentic AI | ZDNET · Technology & Consumer Electronics ^[47]

AGNTCon + MCPCon Japan to Convene Builders Advancing Production-Ready Agentic AI · Technology & Consumer Electronics ^[48]

AGNTCon + MCPCon Japan to Convene Builders Advancing Production-Ready Agentic AI · Technology & Consumer Electronics ^[49]

Inside Healthcare Recruiting's Agentic AI Playbook: Live Panel with HR Executives from ... · Healthcare ^[50]

ai platform AgentSea transforms public healthcare in Singapore | The Straits Times · Healthcare ^[51]

The SWOT of AI in banking - Chris Skinner's blog · Banking, Financial Services & Insurance ^[52]

Closing the Execution Gap: How Agentic AI Drives Faster Supply Chain Decisions · Retail & Ecommerce ^[53]

Granite 4.2 brings native reasoning to enterprise agents - IBM Research · Technology & Consumer Electronics ^[54]

Apple introduces M6 and M5 Ultra for a big leap in performance and AI compute · Technology & Consumer Electronics ^[55]

Broadcom Vs. Micron: Why They're So Strong as Capital Shifts to Inference and Agentic AI · Technology & Consumer Electronics ^[56]

OutSystems Ranks #1 in G2's Enterprise Grid(r) Report for AI App Builders, Named a Leader ... · Technology & Consumer Electronics ^[57]

How do optical networks move beyond cruise control to agentic operations? | Nokia.com · Technology & Consumer Electronics ^[58]

Industrial workforce capacity gap being filled by agentic digital workers - PR Newswire · Energy & Utilities ^[59]

Factor Collaborates with Google Cloud to Help Accelerate Legal AI Transformation with ... · Technology & Consumer Electronics ^[60]

Agentic AI for Retail Stockouts and Returns - Logistics Business(r) Magazine · Retail & Ecommerce ^[61]

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] How, when and why to use agentic AI in our neuroscience labs | The Transmitter · thetransmitter.org
- [2] How to Work with AI Coding Agents | Towards Data Science · towardsdatascience.com
- [3] Beyond AI Coding Assistants: Enterprise-Ready Software Delivery - Happiest Minds · happiestminds.com
- [4] Agentic AI Success Relies On Excellent Human Scaffolding - Semiconductor Engineering · semiengineering.com
- [5] The Real Risk of AI Agents Isn't the Model, It's Governance - Mexico Business News · mexicobusiness.news
- [6] When an AI Agent Acts for the Company, Who Is Liable? - Mexico Business News · mexicobusiness.news
- [7] Agentic AI Contracts Expose Organizations to Runaway Costs With Limited Recourse, Says ... · prnewswire.com
- [8] Sapiens files trademark for SapiensAI - Coverager · coverager.com
- [9] Zywave 2026 Midyear Market Outlooks: Agentic AI a Defining Force in Shaping Coverage ... · prnewswire.com
- [10] The foundations Asia Pacific life insurers need to scale agentic AI with confidence - Deloitte · deloitte.com
- [11] nCino's Q2 FY2027: Agentic AI Banking Thesis Meets Margin Reality - The Futurum Group · futurumgroup.com
- [12] Okta's Early Traction in Agentic AI 'Encouraging,' RBC Says - Yahoo Finance · finance.yahoo.com
- [13] Following an agentic cull, rising adoption is delivering ROI for Workday customers, says ... · diginomica.com
- [14] 5 big analyst AI moves: Downgrades for SAP and Intuit; AMD lifted to Strong Buy · uk.finance.yahoo.com
- [15] Best Buy Advances in Agentic AI After Solid Q2 Earnings, Comps - HomePage News · homepagenews.com
- [16] Intel Outlines Architectures for Agentic AI at Hot Chips 2026 · newsroom.intel.com
- [17] Diamonds, Crescents and Wildcats: Intel shows off its hardware for the next generation of ... · techradar.com
- [18] Intel Crescent Island 160GB to 480GB LPDDR5X AI GPU at Hot Chips 2026 · servethehome.com
- [19] SpaceXAI Adopts NVIDIA Vera CPU to Accelerate Agentic AI at Massive Scale · nvidianews.nvidia.com
- [20] Solving Agentic AI Fleet Challenges with NVIDIA Vera CPU · developer.nvidia.com
- [21] NVIDIA Groq 3 LPX Now in Full Production With World-Class Speed for Agentic AI · markets.businessinsider.com

- [22] [NVIDIA Vera Rubin NVL72 Sets a New Efficiency Standard for AI Agents](#) · [blogs.nvidia.com](#)
- [23] [AWS and NVIDIA to Deliver 2 Million Additional GPUs and Next-Generation Infrastructure for ...](#) · [nvidianews.nvidia.com](#)
- [24] [Matt Garman: The 100 Most Influential People in AI 2026](#) - TIME · [time.com](#)
- [25] [Amazon's AI Investments Are Creating a Whole New Business Model - 24/7 Wall St.](#) · [247wallst.com](#)
- [26] [Alibaba Cloud launches first cloud region in Brazil with agentic AI services](#) · [developingtelecoms.com](#)
- [27] [New Chairman Joel Meyer Dishes On Prebid's Reset For The Agentic Ad Tech Era](#) · [adexchanger.com](#)
- [28] [When Robots Buy Media: Future Maps the Agentic AI Shift in Programmatic](#) - Beet.TV · [beet.tv](#)
- [29] [Agentic AI will Put Media Budgets In Motion: Dentsu X's Brad Stockton](#) - Beet.TV · [beet.tv](#)
- [30] [The future of agentic AI will not be prompt-based | The Current](#) · [thecurrent.com](#)
- [31] [UAT-10147 Uses Agentic AI for Post-Compromise Attacks](#) - SOC Prime · [socprime.com](#)
- [32] [When AI infrastructure becomes the target: Securing gateways and control points](#) - Microsoft · [microsoft.com](#)
- [33] [Agentic AI may make cyberattacks faster, more sophisticated: Securonix CEO](#) · [thehindubusinessline.com](#)
- [34] [Why is AI going on a hacking spree?](#) - Live Science · [livescience.com](#)
- [35] [TrendAI\(tm\) Ranks First on CyberGym Agentic AI Security Benchmark](#) - Asia News Network · [asianews.network](#)
- [36] [An adaptive agentic AI framework for counterfactual validation and adversarial lookahead in ...](#) · [link.springer.com](#)
- [37] [Imagine the SOC Without a Queue: From Alert Backlog to AI Hypothesis Engine](#) · [thehackernews.com](#)
- [38] [Ambient Scribes Improve Clinician Retention, Cleveland Clinic Research Shows](#) · [medcitynews.com](#)
- [39] [The Frontier of Patient Communication: Market Dynamics, Technical Architectures and ...](#) · [ramaonhealthcare.com](#)
- [40] [AgentX - InferenceXv3: Does CUDA Moat Hold up in Agentic Inferencing?](#) - SemiAnalysis · [newsletter.semanalysis.com](#)
- [41] [Is it legal to train AI models on copyrighted books? It's complicated | TechCrunch](#) · [techcrunch.com](#)
- [42] [Google just put agentic hotel booking into testing](#) - Hospitality Net · [hospitalitynet.org](#)
- [43] [AI Era Corp. Completes Its Agentic AI Production Loop](#) - Yahoo Finance Singapore · [sg.finance.yahoo.com](#)
- [44] [EverestLabs Launches First-Ever Agentic AI Platform For Materials Processing, Recovery ...](#) · [prweb.com](#)
- [45] [AI Supercharges New Era of Test & Measurement](#) - IEEE Spectrum · [spectrum.ieee.org](#)
- [46] [EverestLabs Launches First-Ever Agentic AI Platform](#) - Waste 360 · [waste360.com](#)
- [47] [How to use ChatGPT Work - and my top 10 tips for getting started with agentic AI | ZDNET](#) · [zdnet.com](#)
- [48] [AGNTCon + MCPCon Japan to Convene Builders Advancing Production-Ready Agentic AI](#) · [prnewswire.com](#)
- [49] [AGNTCon + MCPCon Japan to Convene Builders Advancing Production-Ready Agentic AI](#) · [sg.finance.yahoo.com](#)
- [50] [Inside Healthcare Recruiting's Agentic AI Playbook: Live Panel with HR Executives from ...](#) · [shrm.org](#)
- [51] [ai platform AgentSea transforms public healthcare in Singapore | The Straits Times](#) · [straitstimes.com](#)
- [52] [The SWOT of AI in banking](#) - Chris Skinner's blog · [thefinanser.com](#)
- [53] [Closing the Execution Gap: How Agentic AI Drives Faster Supply Chain Decisions](#) · [scmr.com](#)
- [54] [Granite 4.2 brings native reasoning to enterprise agents](#) - IBM Research · [research.ibm.com](#)
- [55] [Apple introduces M6 and M5 Ultra for a big leap in performance and AI compute](#) · [apple.com](#)
- [56] [Broadcom Vs. Micron: Why They're So Strong as Capital Shifts to Inference and Agentic AI](#) · [247wallst.com](#)
- [57] [OutSystems Ranks #1 in G2's Enterprise Grid\(r\) Report for AI App Builders, Named a Leader ...](#) · [prnewswire.com](#)
- [58] [How do optical networks move beyond cruise control to agentic operations?](#) | [Nokia.com](#) · [nokia.com](#)
- [59] [Industrial workforce capacity gap being filled by agentic digital workers](#) - PR Newswire · [prnewswire.com](#)
- [60] [Factor Collaborates with Google Cloud to Help Accelerate Legal AI Transformation with ...](#) · [googlecloudpresscorner.com](#)
- [61] [Agentic AI for Retail Stockouts and Returns](#) - Logistics Business(r) Magazine · [logisticsbusiness.com](#)