LUMO INSIGHTS

Thursday, 21 May 2026

From chatbots to co-workers: autonomous agents change the game

Executive Summary

Over the past 48 hours, we've seen major strides in autonomous AI agents - from tech giants launching always-on AI that can act across their ecosystems, to new enterprise tools that use agents to handle complex workflows. These developments show AI moving beyond chatbots and into real business processes, delivering surprising efficiency gains in industries from retail to finance. At the same time, regulators and experts are raising fresh concerns and guidance to ensure these powerful systems remain secure and accountable.

Big tech bets on always-on agents

At Google's annual I/O conference this week, CEO Sundar Pichai declared the dawn of an "Agentic Gemini Era," unveiling the company's most ambitious AI roadmap yet and signalling a shift from AI-enhanced features to fully autonomous AI experiences across the Google ecosystem (economictimes.indiatimes.com [1]). Google announced a barrage of new AI capabilities: a powerful, faster model called Gemini 3.5 "Flash", a multimodal model Gemini Omni that can create content from video to text, an agent-centric development framework called Agentic Antigravity 2.0, and a personal AI assistant named Gemini Spark that runs continuously in the cloud on a user's behalf (economictimes.indiatimes.com [2]) (economictimes.indiatimes.com [3]). This always-on Spark agent can perform tasks in the background - managing emails, scheduling, fetching information, even taking actions in Chrome - without requiring the user to be online (www.eweek.com [4]) (www.eweek.com [5]).

Google's push underscores a larger trend: after years of AI assistants that only advised or summarised, we now have mainstream platforms aiming for autonomous execution. Pichai described these 24/7 background AIs as "personalised AI agents you can set up to work in the background, 24/7, to find what you need at exactly the right moment, and help you take action" (www.eweek.com [6]). In other words, Google is moving beyond answering questions to launching full-fledged digital colleagues that can act on their own. The scale of this transformation is enormous: Google reports its AI systems now process over 3.2 quadrillion tokens per month, a sevenfold increase from a year ago (www.eweek.com [7]), with more than 8.5 million developers actively building with its Gemini models (www.eweek.com [8]).

Other tech leaders are heading in the same direction. Microsoft recently repositioned its Office Copilot as an "Agentic Work" layer, enabling it to autonomously execute multi-step tasks across Outlook, Teams, and Excel with built-in checkpoints for human approval on sensitive actions (advancedai.com [9]) (advancedai.com [10]). Salesforce has similarly overhauled its AI platform: what used to be called Einstein is evolving into Agentforce, reflecting a shift from simple predictive suggestions to AI-driven agents that carry out CRM tasks independently (techradar.info [11]) (techradar.info [12]). The

message is clear: big tech is betting that the next leap in productivity will come from AI that doesn't just assist humans, but works alongside them as autonomous digital agents.

References:
[1] economictimes.indiatimes.com — https://economictimes.indiatimes.com/magazines/panache/google-i/o-2026-google-enters-its-agentic-gemini-era/articleshow/131208621.cms#:~:text=Synopsis%20At%20Google%20I%2FO%202026%2C,signaling%20a%20shift%20from%20AI
[2] economictimes.indiatimes.com — https://economictimes.indiatimes.com/magazines/panache/google-i/o-2026-google-enters-its-agentic-gemini-era/articleshow/131208621.cms#:~:text=YouTube%2C%20Docs%20Live%2C%20Gemini%20Omni%2C,The%20keynote%20focused%20on
[3] economictimes.indiatimes.com — https://economictimes.indiatimes.com/magazines/panache/google-i/o-2026-google-enters-its-agentic-gemini-era/articleshow/131208621.cms#:~:text=workloads%20to%20Gemini%203,building%20and%20orchestrating%20autonomous%20AI
[4] www.eweek.com — https://www.eweek.com/news/google-io-gemini-agentic-ai-era-2026/#:~:text=dashboards%20for%20long,AI%20Ultra%20subscribers%20in%20the
[5] www.eweek.com — https://www.eweek.com/news/google-io-gemini-agentic-ai-era-2026/#:~:text=agent%20inside%20the%20Gemini%20app,%E2%80%9CIt%E2%80%99s%2024%2F7%20so%20you%20don%E2%80%99t
[6] www.eweek.com — https://www.eweek.com/news/google-io-gemini-agentic-ai-era-2026/#:~:text=useful%20updates%20or%20actions%20in,machines%20within%20Google%20Cloud%20infrastructure
[7] www.eweek.com — https://www.eweek.com/news/google-io-gemini-agentic-ai-era-2026/#:~:text=to%20the%20company%2C%20monthly%20token,last%20year%E2%80%99s%20figure%20of%20400
[8] www.eweek.com — https://www.eweek.com/news/google-io-gemini-agentic-ai-era-2026/#:~:text=than%203,last%20year%E2%80%99s%20figure%20of%20400
[9] advancedai.com — https://advancedai.com/briefings/microsoft-copilot-agentic-2026/#:~:text=is%20repositioning%20Copilot%20as%20an,%E2%80%94%20read%20an%20email%20in
[10] advancedai.com — https://advancedai.com/briefings/microsoft-copilot-agentic-2026/#:~:text=Copilot%20Is%20an%20%E2%80%9CAgent%E2%80%9D%3F%20When,doing%20things%20you%20haven%E2%80%99t%20pre
[11] techradar.info — https://techradar.info/salesforce-einstein-is-now-agentforce-the-2026-ai-rebrand-guide/#:~:text=2026%20Salesforce%20Einstein%20is%20Now,driven%20customer%20engagement
[12] techradar.info — https://techradar.info/salesforce-einstein-is-now-agentforce-the-2026-ai-rebrand-guide/#:~:text=Einstein%20served%20as%20the%20umbrella,digital%20labor%20force%20rather%20than

Autonomous agents at the customer frontline

The customer-facing side of business is rapidly being reshaped by AI. Google's new Universal Commerce Protocol (UCP) and related features, announced May 20, establish an ecosystem where AI-driven assistants can handle a consumer's entire buying journey (blog.google [1]). A shopper might ask Google's AI to find a product, then see personalised Gemini-powered conversational ads in Search that offer recommendations with transparent explanations (blog.google [2]), and even have an AI "Universal Cart" assemble items from multiple stores and automatically check out via Google Pay (blog.google [3]). For product and marketing teams, this is a practical shift from agent-assisted discovery to agent-completed purchases, as AI can fill carts, surface promotions, and finalise transactions autonomously - potentially shortening the path from search to sale.

This shift towards agentic commerce could compress sales cycles and alter how companies engage customers. Consumers may soon expect AI "shopping concierge" experiences: an agent that can compare options, answer detailed questions, apply promotions, and execute purchases in a single conversation. Businesses will need to ensure their product data and pricing are easily accessible to such digital agents, and adapt marketing strategies for new formats like chat-based discovery and AI-driven shopping ads (blog.google [4]). Early retail adopters in Google's pilot are already selling directly through AI-curated results, which could boost conversion rates but may also disrupt traditional e-commerce channels.

Other industries are following suit. In financial services, a Temasek-backed startup called Motif this week launched an AI platform named Clarity that provides wealth managers with "structured, sourced" market insights beyond what even dedicated analyst teams have produced (www.cityam.com [5]). By ingesting only verified, high-quality data - from regulatory filings to earnings reports - Clarity's agent can analyse how markets and assets interconnect and change over time, aiming to eliminate AI "hallucinations" in financial research. Multiple institutions had already signed on ahead of launch, collectively representing over 1.5 million end-users and billions in assets under

management (www.cityam.com [6]). Whether in shopping or investing, the common thread is that AI agents are moving into front-line roles, guiding customer decisions and directly driving revenue.

References:
[1] blog.google — https://blog.google/products-and-platforms/products/shopping/shopping-updates-google-marketing-live/#:~:text=generated%20summary%20Google%20is%20launching,options%20make%20it%20easier%20for
[2] blog.google — https://blog.google/products/ads-commerce/google-marketing-live-search-ads/#:~:text=experiences%20that%20offer%20product%20guidance,New%20shopping%20features
[3] blog.google — https://blog.google/products-and-platforms/products/shopping/shopping-updates-google-marketing-live/#:~:text=generated%20summary%20Google%20is%20launching,options%20make%20it%20easier%20for
[4] blog.google — https://blog.google/products/ads-commerce/google-marketing-live-search-ads/#:~:text=experiences%20that%20offer%20product%20guidance,New%20shopping%20features
[5] www.cityam.com — https://www.cityam.com/motif-launches-clarity-a-first-of-its-kind-ai-financial-intelligence-system/#:~:text=connections%20that%20deliver%20the%20kind,com%2Fnews%2Fhome%2F20260511241653%2Fen%2F%20Financial%20institutions
[6] www.cityam.com — https://www.cityam.com/motif-launches-clarity-a-first-of-its-kind-ai-financial-intelligence-system/#:~:text=shifts%20and%20plan%20accordingly,from%20analyst%20research%20and%20regulatory

Autonomous operations and supply chains

AI autonomy is also emerging in the less glamorous but critical realms of supply chain and operations. This week, supply chain software leader **Blue Yonder** announced a partnership with NVIDIA to create a “Model Training Factory” for developing specialised logistics AI agents (diginomica.com [1]). Blue Yonder’s CEO, Duncan Angove, argues that companies should own domain-specific AI “rather than [depend] on costly and imprecise frontier models” (diginomica.com [2]). By fine-tuning open-source models on company data via NVIDIA’s new Nemotron toolset, Blue Yonder aims to deliver supply-chain agents that are more reliable, faster, and far cheaper to run at scale than large general-purpose AIs. For enterprises in retail and manufacturing, this strategy - investing in “owned” AI models for core operations - promises better control, lower latency, and long-term cost advantages.

In workforce management, AI is already tackling complex scheduling challenges. On May 19, HR tech firm **Humanforce** launched an AI-driven **Smart Scheduling** assistant to optimise staffing for frontline teams (www.prnewswire.com [3]). The system automatically aligns shift rosters with predicted demand, reducing managers’ time spent on scheduling by up to 70% and helping cut labour costs by 15% through more efficient staffing (www.prnewswire.com [4]). This kind of efficiency has immediate appeal in sectors like healthcare, retail, and hospitality where staffing is both critical and costly. By taking over the tedium of building schedules or routing deliveries, well-designed agents free managers to focus on exceptions and improvements instead of constant fire-fighting.

The benefits are not just theoretical. One recent survey found that live AI agent projects in back-office functions deliver a median 171% return on investment within 12 months (bananalabs.io [5]). Early adopters of autonomous workflows report significant cost savings, faster cycle times, and the ability to scale operations without linear increases in headcount. As agents handle high-volume repetitive work, human teams can concentrate on higher-value activities, making the organisation more agile and resilient.

References:
[1] diginomica.com — https://diginomica.com/icon-2026-blue-yonder-partners-nvidia-its-model-factory-it-targets-owned-not-rented-intelligence#:~:text=it%20is%20calling%20a%20Model,from%20the%20ground%20up
[2] diginomica.com — https://diginomica.com/icon-2026-blue-yonder-partners-nvidia-its-model-factory-it-targets-owned-not-rented-intelligence#:~:text=intelligence%20By%20Derek%20du%20Preez,Image%20taken%20by
[3] www.prnewswire.com — https://www.prnewswire.com/news-releases/humanforce-launches-ai-powered-smart-scheduling-to-transform-frontline-workforce-management-302775468.html#:~:text=aligns%20labor%20to%20forecasted%20demand,Smart%20Scheduling%20turns
[4] www.prnewswire.com — https://www.prnewswire.com/news-releases/humanforce-launches-ai-powered-smart-scheduling-to-transform-frontline-workforce-management-302775468.html#:~:text=aligns%20labor%20to%20forecasted%20demand,Smart%20Scheduling%20turns
[5] bananalabs.io — https://bananalabs.io/blog/ai-agent-roi#:~:text=Production%20AI%20agents%20delivered%20a,bottom%20quartile%20delivered%20negative%20ROI

Reimagining workflows for an AI-First world

As AI capabilities advance, leaders are realising that automating tasks piecemeal isn't enough - entire processes may need to be redesigned. At CamundaCon this week, enterprise software firm **Camunda** unveiled **ProcessOS**, calling it an "agentic operating system for AI-first enterprise transformation" (www.businesswire.com [1]). Camunda's CEO Jakob Freund argued that the main barrier now is no longer access to AI itself, but determining *where* and *how* to apply it by rebuilding operations around outcomes. "Every process in your enterprise is legacy - it was designed for a world where AI did not exist. This is why we are now entering the decade of 'the great re-engineering': every company will reinvent itself or die" (www.businesswire.com [2]).

Early adopters are beginning to put this into practice. Global bank **Barclays** is among the first to pilot ProcessOS in a closed beta, applying it to complex workflows such as claims processing and customer onboarding. The system uses AI to mine existing logs and documents for current process flows, then automatically generates improved workflows optimised for an AI-driven environment (www.businesswire.com [3]). ProcessOS can incorporate a library of AI agents, data integrations, and human approval checkpoints into these designs, enabling companies to deploy new processes in days instead of months. In one example, a manager can simply describe a desired outcome in natural language and let the platform produce a compliant, optimised process ready to run, with AI agents handling the repetitive tasks and people focusing on exceptions (www.businesswire.com [4]).

This deeper shift is giving rise to new roles and responsibilities. Forward-looking firms are creating job titles like "Agent Supervisor", "Agent QA Lead", "AI Ops Manager", and even **Chief AI Officer** to oversee autonomous systems (www.salesforce.com [5]). These leaders manage how AI agents are trained, governed, and integrated - a critical function as AI becomes central to everyday operations. In one survey, 90% of tech leaders said AI agents are already changing how their teams work, refocusing employees on more strategic, creative tasks and reducing time spent on routine execution (claude.com [6]). The upshot: truly benefiting from agent-driven automation will require rethinking not just technology, but also organisational structures and processes to fully weave AI into the fabric of the enterprise.

References:

[1] www.businesswire.com — https://www.businesswire.com/news/home/20260520352437/en/Camunda-announces-ProcessOS-an-agentic-operating-system-for-AI-first-enterprise-transformation#:~:text=Camunda%20announces%20ProcessOS%2C%20an%20agentic,orchestration%20conference%2C%20in%20front%20of

[2] www.businesswire.com — https://www.businesswire.com/news/home/20260520352437/en/Camunda-announces-ProcessOS-an-agentic-operating-system-for-AI-first-enterprise-transformation#:~:text=outcomes%2C%20not%20forward%20from%20current,Freund%2C%20CEO%2C%20Camunda%20Introducing%20ProcessOS

[3] www.businesswire.com — https://www.businesswire.com/news/home/20260520352437/en/Camunda-announces-ProcessOS-an-agentic-operating-system-for-AI-first-enterprise-transformation#:~:text=agentic%20orchestration%20platform%2C%20which%20processes,growing%20catalog%20of%20extensions%20on

[4] www.businesswire.com — https://www.businesswire.com/news/home/20260520352437/en/Camunda-announces-ProcessOS-an-agentic-operating-system-for-AI-first-enterprise-transformation#:~:text=is%20coming%20for%20business%20operations,CTO%2C%20Camunda%20ProcessOS%20maintains%20a

[5] www.salesforce.com — https://www.salesforce.com/blog/ai-agent-trends-2026/#:~:text=include%20job%20titles%20like%20Agent,more%20AI%20news%20and%20insights

[6] claude.com — https://claude.com/blog/how-enterprises-are-building-ai-agents-in-2026#:~:text=existing%20systems%20%2846,models%20optimized%20for%20coding%20and

Governing the rise of autonomous Work

As AI agents become more powerful, regulators and security leaders are sounding alarms about new risks. A report released by NIST on May 18 synthesised industry feedback on autonomous AI systems, concluding that these agents present novel security threats and that existing cybersecurity practices "require adaptation" for the agent era (www.nist.gov [1]). The U.S. and its Five Eyes allies have likewise urged organisations to introduce AI agents carefully and with strong oversight

(www.theregister.com [2]) (www.theregister.com [3]). Without proper controls, giving software broad autonomy could “amplif[y] organisations’ existing frailties” and open up new, interconnected attack surfaces (www.theregister.com [4]) (www.theregister.com [5]). In one government scenario, a malicious insider asks an AI assistant with system access to “apply a security patch to all servers and clean up the firewall logs,” leading the overly obedient agent to delete critical security records along with the updates (www.theregister.com [6]). The lesson: even as they automate tasks, autonomous agents still need stringent guardrails.

National security agencies have outlined what some of those guardrails should include, identifying five distinct risk categories for agentic AI: *privilege* (ensuring minimal access rights), *design and configuration* flaws, *behavioural* risks (like goal misalignment or deception), *structural* vulnerabilities from integrated components and external data, and *accountability* gaps (lack of oversight and auditability) (startupfortune.com [7]). To address these risks, the joint guidance recommends measures such as giving each agent a verified, cryptographically secure identity, enforcing least-privilege access, implementing human approval for high-impact actions, and guarding against malicious prompt injections (startupfortune.com [8]). In essence, organisations should treat any autonomous AI system as a highly privileged digital employee - one that demands rigorous identity controls, continuous monitoring, and clearly defined limits on its authority.

Industry players are also stepping up with new safety tools. This week NVIDIA unveiled a “Verified Agent Skills” programme, providing a pipeline to audit and digitally sign the software plugins that AI agents use (developer.nvidia.com [9]). The goal is to ensure an agent only executes trusted, well-vetted code - reducing the risk of bad or compromised tools causing harm. And to help enterprises keep sensitive data in-house, Anthropic introduced a self-hosted Claude “sandbox” that allows its AI agents to operate within a company’s own infrastructure (claude.com [10]). Through an encrypted tunnel connection, Claude can now retrieve internal information or trigger software tools without that data ever leaving the corporate firewall (www.infoq.com [11]).

Looking ahead, we can expect more official standards and regulations for autonomous AI. NIST’s initial findings point to principles like resilience (agents must fail safely), reversibility (having “undo” mechanisms for agent actions), robust provenance tracking, and better information-sharing on incidents. Forward-thinking leaders are already adopting such practices - including zero-trust security and rigorous audit trails for AI-driven processes (cyberscoop.com [12]) - as they begin to scale up autonomous workflows. By baking governance into their AI initiatives from the start, enterprises can reap the efficiency benefits of autonomy while staying ahead of compliance requirements and avoiding unwelcome surprises.

References:

[1] www.nist.gov — https://www.nist.gov/publications/summary-analysis-responses-request-information-regarding-security-considerations-ai#:~:text=government%20could%20offer%20to%20facilitate,Commenters

[2] www.theregister.com — https://www.theregister.com/security/2026/05/04/five-eyes-warn-agentic-ai-is-too-dangerous-for-rapid-rollout/5229103#:~:text=Canada%20Simon%20Sharwood%2015%20Published,it%20%E2%80%9Ccrucial%20for%20defenders%20to

[3] www.theregister.com — https://www.theregister.com/security/2026/05/04/five-eyes-warn-agentic-ai-is-too-dangerous-for-rapid-rollout/5229103#:~:text=infrastructure%20from%20agentic%20AI,agent%20empowered%20to%20install%20software

[4] www.theregister.com — https://www.theregister.com/security/2026/05/04/five-eyes-warn-agentic-ai-is-too-dangerous-for-rapid-rollout/5229103#:~:text=Canada%20Simon%20Sharwood%2015%20Published,AI

[5] www.theregister.com — https://www.theregister.com/security/2026/05/04/five-eyes-warn-agentic-ai-is-too-dangerous-for-rapid-rollout/5229103#:~:text=agentic%20AI%20systems%20may%20behave,in%20an%20agentic%20AI%20system

[6] www.theregister.com — https://www.theregister.com/security/2026/05/04/five-eyes-warn-agentic-ai-is-too-dangerous-for-rapid-rollout/5229103#:~:text=document%20offers%20the%20example%20of,AI%20to%20autonomously%20manage%20procurement

[7] startupfortune.com — https://startupfortune.com/six-intelligence-agencies-just-told-enterprise-builders-that-agentic-ai-is-a-live-security-risk-and-the-guidance-is-more-specific-than-anyone-expected/#:~:text=the%20specificity%20of%20its%20recommendations,treats%20the%20agent%20as%20a

[8] startupfortune.com — https://startupfortune.com/six-intelligence-agencies-just-told-enterprise-builders-that-agentic-ai-is-a-live-security-risk-and-the-guidance-is-more-specific-than-anyone-expected/#:~:text=Canada%2C%20and%20New%20Zealand%20published,injection%20defenses%2C%20and%20audit

[9] developer.nvidia.com — https://developer.nvidia.com/blog/nvidia-verified-agent-skills-provide-capability-governance-for-ai-agents/#:~:text=instruction%20sets%20that%20help%20developers,authenticity%20and%20integrity%20beyond%20implied
[10] claude.com — https://claude.com/blog/claude-managed-agents-updates#:~:text=preview%20%28request%20access%29,image%20generation%20get%20the%20CPU
[11] www.infoq.com — https://www.infoq.com/news/2026/05/claude-mcp-tunnels/#:~:text=internal%20systems%20to%20leave%20their,allowing%20organizations%20to%20manage%20compute
[12] cyberscoop.com — https://cyberscoop.com/cisa-nsa-five-eyes-guidance-secure-deployment-ai-agents/#:~:text=match%20at%20L67%20guidance%20urges,unmonitored%20access%20to%20sensitive%20networks

Key Statistics

- 3.2 quadrillion - monthly tokens processed by Google's AI products, a 7x increase from last year ([www.eweek.com](https://www.eweek.com/news/google-io-gemini-agentic-ai-era-2026/#:~:text=to%20the%20company%2C%20monthly%20token,last%20year%E2%80%99s%20figure%20of%20400))
- 8.5 million - developers building applications with Google's Gemini models each month ([www.eweek.com](https://www.eweek.com/news/google-io-gemini-agentic-ai-era-2026/#:~:text=than%203,last%20year%E2%80%99s%20figure%20of%20400))
- 80% - share of organisations reporting measurable ROI from their AI agent investments ([claude.com](https://claude.com/blog/how-enterprises-are-building-ai-agents-in-2026#:~:text=Looking%20ahead%2C%2056,of%20case%20law%20and%203%2C000))
- 171% - median one-year ROI for companies that put AI agents into production (2026 survey) ([bananalabs.io](https://bananalabs.io/blog/ai-agent-roi#:~:text=Production%20AI%20agents%20delivered%20a,bottom%20quartile%20delivered%20negative%20ROI))
- 70% - reduction in time required to create staff rosters using an AI scheduling assistant (with ~15% lower labour costs) ([www.prnewswire.com](https://www.prnewswire.com/news-releases/humanforce-launches-ai-powered-smart-scheduling-to-transform-frontline-workforce-management-302775468.html#:~:text=aligns%20labor%20to%20forecasted%20demand,Smart%20Scheduling%20turns))
- 5 - distinct risk categories for autonomous AI identified by joint Five Eyes security guidance ([startupfortune.com](https://startupfortune.com/six-intelligence-agencies-just-told-enterprise-builders-that-agentic-ai-is-a-live-security-risk-and-the-guidance-is-more-specific-than-anyone-expected/#:~:text=the%20specificity%20of%20its%20recommendations,treats%20the%20agent%20as%20a))

KEY TAKEAWAY

Leaders must seize these advances now by launching controlled pilots of AI agents in high-impact workflows; in parallel, they should redesign processes and governance to safely scale autonomy across the enterprise

Sources

How enterprises are building AI agents in 2026 (Claude blog)
https://claude.com/blog/how-enterprises-are-building-ai-agents-in-2026

Google I/O 2026: Google Enters Its 'Agentic Gemini Era' - The Economic Times
https://economictimes.indiatimes.com/magazines/panache/google-i/o-2026-google-enters-its-agentic-gemini-era/articleshow/131208621.cms

New ad formats built with Gemini coming to Google Search (Google Blog)
https://blog.google/products/ads-commerce/google-marketing-live-search-ads/

How we're helping retailers thrive with new UCP and AI tools (Google Blog)
https://blog.google/products-and-platforms/products/shopping/shopping-updates-google-marketing-live/

Google I/O 2026: All the Major AI Announcements - eWeek
https://www.eweek.com/news/google-io-gemini-agentic-ai-era-2026/

Microsoft Copilot Is Now an Agent — What That Means for Your M365 Team - Advanced AI (April 2026)
https://advancedai.com/briefings/microsoft-copilot-agentic-2026/

Salesforce Einstein is Now Agentforce: The 2026 AI Rebrand Guide - TechRadar.info
https://techradar.info/salesforce-einstein-is-now-agentforce-the-2026-ai-rebrand-guide/

Camunda announces ProcessOS, an agentic operating system for AI-first enterprise transformation - Business Wire
https://www.businesswire.com/news/home/20260520352437/en/Camunda-announces-ProcessOS-an-agentic-operating-system-for-AI-first-enterprise-transformation

ICON 2026: Blue Yonder partners with NVIDIA for its model factory, targeting “owned” supply chain AI - Diginomica
https://diginomica.com/icon-2026-blue-yonder-partners-nvidia-its-model-factory-it-targets-owned-not-rented-intelligence

Humanforce launches AI-powered Smart Scheduling to transform frontline workforce management - PR Newswire
https://www.prnewswire.com/news-releases/humanforce-launches-ai-powered-smart-scheduling-to-transform-frontline-workforce-management-301830366.html

New in Claude Managed Agents: self-hosted sandboxes and MCP tunnels (Anthropic blog)
https://claude.com/blog/claude-managed-agents-updates

Anthropic Introduces MCP Tunnels for Private Agent Access to Internal Systems - InfoQ
https://www.infoq.com/news/2026/05/claude-mcp-tunnels/

Five Eyes warn agentic AI is too dangerous for rapid rollout - The Register
https://www.theregister.com/security/2026/05/04/five-eyes-warn-agentic-ai-is-too-dangerous-for-rapid-rollout/5229103

US government, allies publish guidance on how to safely deploy AI agents - CyberScoop
https://cyberscoop.com/cisa-nsa-five-eyes-guidance-secure-deployment-ai-agents/

Six Intelligence Agencies Just Told Enterprise Builders That Agentic AI Is a Live Security Risk... - Startup Fortune
https://startupfortune.com/six-intelligence-agencies-just-told-enterprise-builders-that-agentic-ai-is-a-live-security-risk-and-the-guidance-is-more-specific-than-anyone-expected/

8 Ways AI Agents Are Evolving in 2026 - Salesforce Blog
https://www.salesforce.com/blog/ai-agent-trends-2026/

NVIDIA-Verified Agent Skills Provide Capability Governance for AI Agents - NVIDIA Technical Blog
https://developer.nvidia.com/blog/nvidia-verified-agent-skills-provide-capability-governance-for-ai-agents/

Summary Analysis of Responses to the RFI on Security Considerations for AI Agents - NIST (May 2026)
https://nvlpubs.nist.gov/nistpubs/ai/NIST.IR.8477.pdf