LUMO INSIGHTS

Thursday, 24 September 2026

AI agents: cost breakthroughs, team-ups and new risks in 48 hours

Executive Summary

Over the last two days, leading AI providers launched dramatically cheaper, more capable agent AI models, accelerating enterprise adoption of autonomous workflows. At the same time, new real-world developments - from collaborative AI teams to AI-driven cyberattacks - revealed both breakthrough opportunities and urgent governance challenges.

Frontier AI models get cheaper and more capable

This week saw an unprecedented price war between two leading AI labs, signalling a shift in how quickly cutting-edge AI becomes affordable. Within 90 minutes on Tuesday, Anthropic unveiled its new Claude Opus 5.5 model and OpenAI responded by launching GPT-6 Sol and Luna - dramatic moves that turned what are usually purely capability-driven frontier model releases into a fight over price. Both companies cut costs sharply: OpenAI's GPT-6 Sol and Luna models are priced at roughly half the per-token rate of the previous generation, undercutting Anthropic's own price drop with Opus 5.5.

(kingy.ai [1])Crucially, these models also deliver technical leaps enabling more ambitious autonomous workflows. For instance, GPT-6 Sol boasts an astounding 1.05 /million-token context window (with outputs up to 128,000 tokens), allowing an AI agent to ingest and reason over vast amounts of information in a single session. Such expanded memory means an agent can analyse entire codebases, years of transaction data, or comprehensive policy documents all at once - tasks that previously had to be broken into smaller chunks due to context limits.

(www.theneuron.ai [2])The capabilities unlocked by these upgrades are already evident. Anthropic reports that Claude Opus 5.5 recently translated a 680,000-line legacy codebase in under one day - a task that would have taken human engineers weeks or months - and conducted a 200,000-line code audit in under three hours. In parallel, OpenAI says GPT-6 Sol can handle complex business process simulations more efficiently than its rivals, completing end-to-end software and support workflows at roughly a tenth of the cost (openai.com [3]). In short, the frontier AI race is now as much about operational efficiency as raw intelligence.

This new cost-performance equation can significantly lower the barrier for enterprises to deploy advanced AI agents at scale. Many tasks that were cost-prohibitive to automate a year ago - from comprehensive customer service bots to AI-assisted financial analysis - may now be economically viable. Business leaders should take a fresh look at their AI pipeline and backlog of projects. Use cases once shelved due to expense or model limitations might now deliver strong ROI, thanks to cheaper tokens and models that can retain vastly more context.

References:

[1] kingy.ai — https://kingy.ai/news/2026-09-22-ai-launch-radar/#:~:text=Why%20it%20matters%3A%20Sol%20adds,a%20measure%20of%20answer%20quality

[2] www.theneuron.ai — https://www.theneuron.ai/digest/everything-that-happened-in-ai-today-tuesday-september-22-2026/#:~:text=picture%20is%20messier,at%20higher%20effort%20it%20matches
[3] openai.com — https://openai.com/index/introducing-gpt-6-sol-and-luna/#:~:text=work%20GPT%E2%80%916%20Sol%20can%20take,30

Enterprise AI strategy shifts to control cost and models

The rush to adopt AI agents is exposing a new strategic challenge: managing the cost and control of the underlying AI models. Reports this week revealed that Harvey - a prominent legal AI startup recently valued at $15.6 billion - saw its gross margins plunge to -50% due to the steep expenses of using OpenAI's and Anthropic's cloud APIs.

(blog.4sapi.com [1]) (blog.4sapi.com [2])In response, Harvey pivoted to develop its own model, adapting an open-source open-weight large model from China (the Kimi K3 from Moonshot AI) as the foundation of its proprietary Harvey Tenet system. This switch reportedly cut Harvey's AI inference costs by roughly 75%, turning an unsustainable cost structure into a viable model. (blog.4sapi.com [3])

Harvey's drastic move underscores an inflection point: as AI agent usage scales, companies must weigh the benefits of best-in-class external models against the long-term economics and dependencies involved. The most advanced services from providers like OpenAI or Anthropic still offer unmatched capabilities, but heavy reliance on them can become prohibitively expensive at enterprise scale. Now that high-performing open models and partner-trained 'co-pilots' are emerging as credible alternatives, even market leaders are rethinking their 'buy vs build' decisions. This may herald a wave of companies developing custom AI models (or adopting open-source ones) tailored to specific domains - a trend already starting to play out in other sectors. (blog.4sapi.com [4])

The lesson for executives is that AI transformation is not a set-and-forget procurement exercise. It demands active cost management and strategic flexibility. Leaders should ensure their AI investments include a plan to continually optimise model selection and hosting strategies. In some cases, renegotiating cloud AI usage contracts, exploring open-source models, or investing in internal model development will be necessary to maintain healthy margins and control over mission-critical AI capabilities. The organisations that thrive with AI will be those that balance innovation with vigilance on business fundamentals.

References:
[1] blog.4sapi.com — https://blog.4sapi.com/blog/harvey-tenet-kimi-k3-legal-ai#:~:text=A%20striking%20shift%20has%20emerged,When%20engineering%20teams
[2] blog.4sapi.com — https://blog.4sapi.com/blog/harvey-tenet-kimi-k3-legal-ai#:~:text=improvement%2C%20doubled%20task%20completion%20rates,enterprise%20teams%20in%20law%2C%20finance
[3] blog.4sapi.com — https://blog.4sapi.com/blog/harvey-tenet-kimi-k3-legal-ai#:~:text=improvement%2C%20doubled%20task%20completion%20rates,enterprise%20teams%20in%20law%2C%20finance
[4] blog.4sapi.com — https://blog.4sapi.com/blog/harvey-tenet-kimi-k3-legal-ai#:~:text=A%20striking%20shift%20has%20emerged,When%20engineering%20teams

From chatbots to AI teams: agents as collaborative workforces

(openclawnews.tech [1])A new class of AI tools is redefining how work gets done by enabling AI agents to collaborate with each other and with humans. This week, OpenClaw 2.0 - an open-source agent platform - demonstrated how a multiplayer agent system can have dozens of autonomous specialist agents sharing one workspace across multiple applications and machines (openclawnews.tech [2]). Unlike a solitary chatbot that simply answers questions, a team of AI agents can pass tasks to one another, persistently maintain context over time, and collectively work towards an overarching goal. Human colleagues can even join these shared sessions to oversee progress or take the reins when needed, then hand tasks back to the agents to continue autonomously.

The promise of such an approach is faster, continuous operations. With persistent context and role-based AI 'colleagues', a project need not pause when an employee signs off - the agent team can keep the momentum going around the clock (openclawnews.tech [3]). For instance, a data analysis that starts in the afternoon could run through the night, with one agent pulling information from various enterprise systems, another performing complex analyses or drafting reports, and a third double-checking for errors - all while a human manager is available at key decision points. Early adopters are taking note: Meta’s new personal AI assistant, Muse, reportedly drew inspiration from OpenClaw’s multi-agent design, resulting in an agent that can schedule meetings, reply to emails, and coordinate tasks across apps as a sort of personal chief of staff (openclawnews.tech [4]).

While multi-agent workflows can dramatically increase productivity, they also raise the bar for governance and IT architecture (openclawnews.tech [5]). As AI systems take on more collective responsibilities, organisations must ensure clarity around data access, permissions, and audit trails within these autonomous teams. A recent security review of OpenClaw’s collaborative agent environment uncovered 23 distinct vulnerabilities - from permission leaks to persistent access bugs (openclawnews.tech [6]) - emphasising that robust identity and access management is as essential for AI agents as it is for human employees. Companies experimenting with agent orchestration platforms should invest early in strong oversight mechanisms (e.g. role-based access controls, real-time monitoring, and clear human-in-the-loop processes) to capture multi-agent efficiency without sacrificing security or accountability.

References:

[1] openclawnews.tech — https://openclawnews.tech/openclaw-2-multiplayer-workforce-security/#:~:text=A%20traditional%20chatbot%20has%20a,system%20has%20more%20moving%20parts
[2] openclawnews.tech — https://openclawnews.tech/openclaw-2-multiplayer-workforce-security/#:~:text=Recent%20coverage%20of%20OpenClaw%E2%80%99s%20development,2
[3] openclawnews.tech — https://openclawnews.tech/openclaw-2-multiplayer-workforce-security/#:~:text=OpenClaw%E2%80%99s%20multiplayer%20direction%20emphasizes%20shared,and%20make%20handoffs%20much%20faster
[4] openclawnews.tech — https://openclawnews.tech/openclaw-2-multiplayer-workforce-security/#:~:text=accomplish%20,2
[5] openclawnews.tech — https://openclawnews.tech/openclaw-2-multiplayer-workforce-security/#:~:text=This%20attention%20is%20arriving%20alongside,authorization%2C%20observability%2C%20and%20recovery%20become
[6] openclawnews.tech — https://openclawnews.tech/openclaw-2-multiplayer-workforce-security/#:~:text=This%20attention%20is%20arriving%20alongside,authorization%2C%20observability%2C%20and%20recovery%20become

Autonomous AI brings new security and trust risks

(blog.rankiteo.com [1])Alongside these advances, recent reports have exposed how AI autonomy can be weaponised or go awry. An analysis of a cybercrime campaign revealed that a threat actor used open-source AI agents to systematically breach 27 companies - including a Fortune 500 hospitality firm and a major airline - and steal over 600,000 customer credit card records (blog.rankiteo.com [2]). Even more alarming, the entire operation ran with almost no human intervention and at an average cost of only $25 per breached site (blog.rankiteo.com [3]). By chaining together multiple publicly available AI tools that can identify software vulnerabilities, craft exploits, and manage intrusions, the attacker massively scaled their reach and adapted in real time to defeat standard defences. It’s a stark illustration that as AI tools become more capable and accessible, the cost of sophisticated attacks falls dramatically - and defenders face a rapidly expanding threat surface.

(www.bitdefender.com [4])Meanwhile, leaders in the financial industry are raising concerns about new risks as AI agents take on more authority in business processes. A consortium of major banks - including Bank of America, Capital One, ING, NatWest, and others - just called for common standards to govern 'agentic commerce,' where AI assistants can execute online purchases and payments for customers (www.bitdefender.com [5]). The banks warn that giving AI power over transactions could create novel avenues for fraud, impersonation, data theft, and compliance failures if proper

safeguards aren’t in place. Among the scenarios they highlight are hacked shopping bots that steal payment information or AI systems over-collecting personal data (www.bitdefender.com [6]). That such heavyweight financial institutions are proactively urging controls shows that even automation-forward sectors are uneasy about the leap from recommendation to autonomous action.

(techiexpert.com [7])Even well-intentioned AI deployments can introduce unintended vulnerabilities if not tightly governed. This week, a popular AI coding assistant was discovered silently uploading entire software codebases - including their version histories - to an external cloud server without users’ knowledge (techiexpert.com [8]). And in healthcare, a new clinical study found that while an AI tool helped oncologists improve lung cancer treatment planning accuracy from 57% to 65%, doctors also occasionally followed the AI’s incorrect suggestions, underscoring the perils of overreliance. These incidents reinforce the need for rigorous oversight, transparency, and human checkpoints as organisations embrace more autonomous AI. To safely reap the benefits of agent-powered automation, companies must redouble efforts on security testing, data governance, staff training to interpret AI output, and clear protocols ensuring human control over critical decisions.

References:

[1] blog.rankiteo.com — https://blog.rankiteo.com/fas1790144700-fashion-retailer-cyber-attack-september-2026/#:~:text=23%2C%202026%20%E2%80%94%203%20min,and%20a%20fashion%20retailer%20between

[2] blog.rankiteo.com — https://blog.rankiteo.com/fas1790144700-fashion-retailer-cyber-attack-september-2026/#:~:text=%2425,with%20minimal%20human%20input%20via

[3] blog.rankiteo.com — https://blog.rankiteo.com/fas1790144700-fashion-retailer-cyber-attack-september-2026/#:~:text=23%2C%202026%20%E2%80%94%203%20min,and%20a%20fashion%20retailer%20between

[4] www.bitdefender.com — https://www.bitdefender.com/en-us/blog/hotforsecurity/ai-agent-online-shopping-banks-scams-privacy#:~:text=agentic%20commerce%2C%20could%20dramatically%20change,When%20your%20AI%20assistant

[5] www.bitdefender.com — https://www.bitdefender.com/en-us/blog/hotforsecurity/ai-agent-online-shopping-banks-scams-privacy#:~:text=agentic%20commerce%2C%20could%20dramatically%20change,When%20your%20AI%20assistant

[6] www.bitdefender.com — https://www.bitdefender.com/en-us/blog/hotforsecurity/ai-agent-online-shopping-banks-scams-privacy#:~:text=cybersecurity%20and%20privacy%20risk,When%20your%20AI%20assistant

[7] techiexpert.com — https://techiexpert.com/zcode-caught-silently-uploading-entire-workspaces-and-git-records-to-the-cloud/#:~:text=,bucket%2C%20and%20brought%20in%20CAICT

[8] techiexpert.com — https://techiexpert.com/zcode-caught-silently-uploading-entire-workspaces-and-git-records-to-the-cloud/#:~:text=,bucket%2C%20and%20brought%20in%20CAICT

Rethinking knowledge as executable agents

(www.theneuron.ai [1])A paradigm-shifting experiment from Stanford this week suggests that AI agents could transform how organisations manage and apply knowledge. Researchers unveiled a system called Paper2Agent that converts a conventional scientific paper into an interactive AI agent capable of running the paper’s core analyses, applying its methods to new data, and even collaborating with other agents built from different papers (www.theneuron.ai [2]).

(www.theneuron.ai [3])Nature’s coverage of this project described the resulting agent as a kind of 'virtual corresponding author' - essentially a software clone of the paper’s original author, able to explain what it knows and perform the work described in the study (www.theneuron.ai [4]). The Paper2Agent system was formally published in the journal Nature, pushing the boundaries of how AI can act as an autonomous researcher and hinting at a future where even a company’s internal knowledge could be turned into living, executable agents instead of static documents.

(www.theneuron.ai [5])While still experimental, this approach offers a glimpse into a future where corporate knowledge isn’t just stored and read, but actively put to work. One analysis noted that the internet made documents searchable and AI made them conversational - now emerging tools like Paper2Agent are beginning to make knowledge executable (www.theneuron.ai [6]). In practical terms, firms might eventually convert their internal reports, training manuals, or research outputs into AI agents that employees can query for instant answers or task with carrying out processes. That could dramatically compress the time from insight to action, provided organisations establish rigorous

validation and oversight to ensure these autonomous knowledge agents remain accurate and aligned with business goals.

References:
[1] www.theneuron.ai — https://www.theneuron.ai/news/stanfords-paper2agent-turns-research-papers-into-working-ai-agents/#:~:text=to%20figure%20out%20how%20to,sounds%20like%20an%20upgraded%20research
[2] www.theneuron.ai — https://www.theneuron.ai/news/stanfords-paper2agent-turns-research-papers-into-working-ai-agents/#:~:text=turn%20it%20into%20an%20interactive,that%20the%20scientific%20paper%20itself
[3] www.theneuron.ai — https://www.theneuron.ai/news/stanfords-paper2agent-turns-research-papers-into-working-ai-agents/#:~:text=and%20even%20collaborating%20with%20agents,just%20ChatGPT%20for%20PDFs%20We%27ve
[4] www.theneuron.ai — https://www.theneuron.ai/news/stanfords-paper2agent-turns-research-papers-into-working-ai-agents/#:~:text=and%20even%20collaborating%20with%20agents,just%20ChatGPT%20for%20PDFs%20We%27ve
[5] www.theneuron.ai — https://www.theneuron.ai/news/stanfords-paper2agent-turns-research-papers-into-working-ai-agents/#:~:text=The%20internet%20made%20scientific%20papers,like%20the%20endpoint%20of%20scientific
[6] www.theneuron.ai — https://www.theneuron.ai/news/stanfords-paper2agent-turns-research-papers-into-working-ai-agents/#:~:text=The%20internet%20made%20scientific%20papers,like%20the%20endpoint%20of%20scientific

Key Statistics

- OpenAI slashed API prices for its new GPT-6 Sol and Luna models by 50%, doubling the token output per dollar ([www.unite.ai](https://www.unite.ai/openai-introduces-gpt-6-sol-and-luna-with-50-lower-api-prices/#:~:text=while%20costing%20dramatically%20less%20per,its%20announcement%2C%20OpenAI%20said%20it)).
- Cybercriminals used AI agents to steal 600,000+ credit card records from 27 companies at an average cost of only ~$25 per target ([blog.rankiteo.com](https://blog.rankiteo.com/fas1790144700-fashion-retailer-cyber-attack-september-2026/#:~:text=23%2C%202026%20%E2%80%94%203%20min,and%20a%20fashion%20retailer%20between)).
- $15.6 billion legal AI startup Harvey built an in-house model on open-source tech and slashed its AI inference costs by ~75% ([blog.4sapi.com](https://blog.4sapi.com/blog/harvey-tenet-kimi-k3-legal-ai#:~:text=improvement%2C%20doubled%20task%20completion%20rates,enterprise%20teams%20in%20law%2C%20finance)).

KEY TAKEAWAY

With cutting-edge AI now far cheaper and more autonomous, leaders should move fast to integrate AI agents into core workflows while ensuring rigorous cost controls, robust security, and clear governance to harness these tools safely.

Sources

OpenAI Introduces GPT-6 Sol and Luna With 50% Lower API Prices
https://www.unite.ai/openai-introduces-gpt-6-sol-and-luna-with-50-lower-api-prices/

Daily AI Launch Radar: September 22, 2026
https://kingy.ai/news/2026-09-22-ai-launch-radar/

Everything That Happened in AI Today (Tuesday, September 22, 2026)
https://www.theneuron.ai/digest/everything-that-happened-in-ai-today-tuesday-september-22-2026/

OpenClaw 2.0's Multiplayer Workforce: Shared Sessions, Remote Compute, and the Security Lessons of 23 Patched Vulnerabilities
https://openclawnews.tech/openclaw-2-multiplayer-workforce-security/

Paper2Agent Turns Scientific Papers Into AI Agents
https://www.theneuron.ai/news/stanfords-paper2agent-turns-research-papers-into-working-ai-agents/

AI agents want to shop for you; banks warn of scams and privacy risks
https://www.bitdefender.com/en-us/blog/hotforsecurity/ai-agent-online-shopping-banks-scams-privacy

AI-Powered Cybercrime Campaign Steals 600K+ Credit Cards from Retailers at $25 per Target
https://cyberpress.org/autonomous-ai-agents-hack-online-retailers-for-25-a-target-steal-600000-credit-cards/

Harvey Tenet: Kimi K3 Cuts Legal AI Costs by 75%
https://blog.4sapi.com/blog/harvey-tenet-kimi-k3-legal-ai

ZCode Caught Silently Uploading Entire Workspaces and Git Records to the Cloud

https://techiexpert.com/zcode-caught-silently-uploading-entire-workspaces-and-git-records-to-the-cloud/