LUMO INSIGHTS

Thursday, 12 March 2026

From support calls to courts: AI agents at work

Executive Summary

In the past 48 hours, AI agents have surged into enterprise workflows. Salesforce's new Agentforce Contact Centre automates roughly 40 - 60% of customer calls (www.itpro.com [1]), Google embedded its Gemini AI into Workspace to auto-generate documents and spreadsheets from corporate data (www.itpro.com [2]) (www.itpro.com [3]), and Nvidia hinted at an open-source agent platform. However, friction has appeared: a US court blocked a third-party AI shopping bot (www.windowscentral.com [4]), and studies caught medical chatbots confidently repeating dangerous misinformation (www.livescience.com [5]).

References:

[1] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=are%20seeing%20over%20half%20of,high%20quality%2C%20without%20a%20subsequent

[2] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/google-workspace-just-got-a-huge-gemini-update-heres-what-to-expect-with-new-ai-features-in-docs-slides-sheets-and-drive#:~:text=Google%20Docs

[3] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/google-workspace-just-got-a-huge-gemini-update-heres-what-to-expect-with-new-ai-features-in-docs-slides-sheets-and-drive#:~:text=With%20spreadsheets%2C%20accuracy%20is%20key%2C,existing%20sheets%20or%20the%20web

[4] www.windowscentral.com — https://www.windowscentral.com/hardware/laptops/amazons-quiet-crackdown-on-ai-shopping-agents-reveals-the-real-reason-youre-not-allowed-to-save-money-online#:~:text=Except%20that%27s%20not%20what%20happened,if%20that%20automation%20is%20named

[5] www.livescience.com — https://www.livescience.com/health/rectal-garlic-insertion-for-immune-support-medical-chatbots-confidently-give-disastrously-misguided-advice-experts-say#:~:text=same%20claim%20was%20repackaged%20in,of%20the%20time

Major platform releases

Tech giants raced to ship new agentic tools. Salesforce's Agentforce platform now ties AI agents directly into its CRM and telephony stack (www.itpro.com [1]). Early demos show an agent taking a customer's frustrated flight-rebooking request, resolving some issues with AI and seamlessly handing off the rest to a human. In one example, the AI recognised the passenger's frustration and transferred him to a live rep - and that rep instantly saw a summary of the interaction so far (www.itpro.com [2]) (www.itpro.com [3]). Microsoft's competing Copilot is similarly moving beyond chat: it now integrates Anthropic's Claude Cowork to automate scheduled tasks (e.g. triaging calendars or briefing docs) and offers a new AI-focused Microsoft 365 tier. Meanwhile, Google quietly rolled out a slew of Gemini upgrades in its Workspace suite. A "Help me create" sidebar in Docs and Slides now lets you describe a content need and have Gemini pull from your Gmail, Drive and web to draft a first version (www.itpro.com [4]). Sheets gained similar auto-build features - for instance, Gemini can read last year's bookkeeping and suggest a formatted budget chart (www.itpro.com [5]). Google reports a ~70% benchmark accuracy (near expert level) on these tasks (www.itpro.com [6]), reinforcing that humans must validate results. These product moves indicate that AI agents are becoming a core part of enterprise app platforms (Salesforce, Microsoft, Google, Nvidia).

Nvidia announced its next gambit: according to leaks, the company is developing "NemoClaw," an open-source platform for building and orchestrating autonomous agents (www.tomshardware.com [7]). Unusually, NemoClaw is said to be hardware-agnostic (not requiring Nvidia GPUs) and pitched to

major partners like Adobe, Google and Salesforce (www.techradar.com [8]). It will bundle developer tooling and security features aimed at workflow automation (imagine chaining tasks across document processing, compliance checks, customer follow-up, etc.). If confirmed at next week's GTC conference, NemoClaw would represent a step-change: Nvidia betting on managing the agent layer, not just selling chips. (www.techradar.com [9]) (www.tomshardware.com [10])

References:

[1] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=Date%3A%20Tue%20Mar%2010%2013%3A19%3A18,2026%20UTC

[2] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=In%20a%20demo%2C%20Gautam%20Vasudev%2C,to%20the%20relevant%20reservation%20details

[3] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=are%20seeing%20over%20half%20of,high%20quality%2C%20without%20a%20subsequent

[4] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/google-workspace-just-got-a-huge-gemini-update-heres-what-to-expect-with-new-ai-features-in-docs-slides-sheets-and-drive#:~:text=Google%20Docs

[5] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/google-workspace-just-got-a-huge-gemini-update-heres-what-to-expect-with-new-ai-features-in-docs-slides-sheets-and-drive#:~:text=With%20spreadsheets%2C%20accuracy%20is%20key%2C,existing%20sheets%20or%20the%20web

[6] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/google-workspace-just-got-a-huge-gemini-update-heres-what-to-expect-with-new-ai-features-in-docs-slides-sheets-and-drive#:~:text=With%20spreadsheets%2C%20accuracy%20is%20key%2C,existing%20sheets%20or%20the%20web

[7] www.tomshardware.com — https://www.tomshardware.com/tech-industry/artificial-intelligence/nvidia-reportedly-building-its-own-ai-agent-to-compete-with-openclaw-report-claims-nemoclaw-will-supposedly-be-open-source-and-designed-for-enterprise-use#:~:text=Nvidia%20is%20reportedly%20planning%20to,will%20work%20on%20any%20hardware

[8] www.techradar.com — https://www.techradar.com/pro/nvidia-might-be-about-to-reimagine-ai-agents-at-work-with-new-nemoclaw-release#:~:text=A%20Wired%20report%20details%20how,source%2C%20which%20could%20see%20contributors

[9] www.techradar.com — https://www.techradar.com/pro/nvidia-might-be-about-to-reimagine-ai-agents-at-work-with-new-nemoclaw-release#:~:text=Despite%20being%20a%20slated%20Nvidia,the%20recent%20popularity%20of%20%27claws

[10] www.tomshardware.com — https://www.tomshardware.com/tech-industry/artificial-intelligence/nvidia-reportedly-building-its-own-ai-agent-to-compete-with-openclaw-report-claims-nemoclaw-will-supposedly-be-open-source-and-designed-for-enterprise-use#:~:text=are%20in%20short%20supply%20,reason%20that%20the%20company%20is

Enterprise agents in action

Real-world deployments are already catching attention. Early adopters report substantial efficiency gains. In travel and hospitality, one company reports that its AI voice agent alone now resolves roughly 40 - 60% of inbound support calls without any human intervention (www.itpro.com [1]). Supervisors see transcripts and sentiment analytics live, allowing them to step in on tricky calls. Nonprofits are also seeing big returns. For example, the nonprofit Compass Working Capital uses Salesforce's agents to automate coach note-taking (capturing barriers, employment status, etc.), and estimates saving about 6,000 staff-hours a year (www.itpro.com [2]). A customer at Smart Home firm Savant says Agentforce helps prioritise which support calls to escalate, freeing her team to focus on complex issues. Through these use-cases, agents are shouldering high-volume, repetitive tasks (like logging case notes or reordering components) while grounding their outputs in the company's own data.

Even personal productivity is being reimagined: in previews, workers can hand off entire preparations to agents. For example, a user told one AI to organise their week's meetings, fetch relevant emails, and draft updated docs - tasks that unused to require a human assistant. Experiments at Microsoft and open-source demos show agents can compile slide decks from text prompts, schedule conferences by scanning calendars, and summarise prior discussion points. The key enabler is integrating agents into everyday tools: when fed CRM records, Slack chats and email threads, these agents can pre-complete tasks like generating a draft project plan or filling in spreadsheets. The message for C-levels is simple: agentic automation isn't hypothetical anymore - it's driving clear outcomes (faster service, fewer manual hours) in live deployments.

References:

[1] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=are%20seeing%20over%20half%20of,high%20quality%2C%20without%20a%20subsequent

[2] www.itpro.com — https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=barriers%20to%20their%20goals%2C%20which,across%20the%20agency%20each%20year

Pushback and surprises

Even as enterprise agents advance, tensions have emerged. In retail e-commerce, Amazon took the lead - and a court - to block third-party shopping bots. Perplexity's 'Comet' browser AI, which could automate purchases for users, was ordered offline because Amazon argued it lacked "authorisation" to use customer accounts (www.windowscentral.com [1]). The court agreed that while users consented, Amazon's terms did not. Critics note this isn't an innocent security move but a business one: Amazon wants to keep you inside its ad-driven ecosystem (www.windowscentral.com [2]). In effect, Amazon says you can have voice-ordering (Alexa) or ads, but not your own deal-finding bot.

Meanwhile, researchers continue to uncover AI agent failures in critical domains. A recent Lancet Digital Health study charged that popular medical chatbots confidently delivered disastrously wrong advice. When false medical claims were phrased in formal clinical language (e.g. recommending "rectal garlic insertion" for immunity), bots failed 46% of the time (www.livescience.com [3]). In contrast, the same bogus advice casually worded was mostly rejectedÿ only ~9% failures) (www.livescience.com [4]). The conclusion: LLMs have learned that "medical-sounding" text is authoritative - so they regurgitate dangerous misinformation when it looks like an expert, even in dissonance with facts. In practice, millions of people already ask ChatGPT-level AI for health tips every day (www.livescience.com [5]). These findings are a stark reminder that unsupervised agents can amplify existing societal blindspots (misinformation, bias) at scale. Enterprises piloting agents in regulated sectors (like healthcare or finance) should not assume proper answers: they must validate outputs as they go.

References:

[1] www.windowscentral.com — https://www.windowscentral.com/hardware/laptops/amazons-quiet-crackdown-on-ai-shopping-agents-reveals-the-real-reason-youre-not-allowed-to-save-money-online#:~:text=Except%20that%27s%20not%20what%20happened,if%20that%20automation%20is%20named
[2] www.windowscentral.com — https://www.windowscentral.com/hardware/laptops/amazons-quiet-crackdown-on-ai-shopping-agents-reveals-the-real-reason-youre-not-allowed-to-save-money-online#:~:text=Amazon%27s%20business%20model%20is%20increasingly,items
[3] www.livescience.com — https://www.livescience.com/health/rectal-garlic-insertion-for-immune-support-medical-chatbots-confidently-give-disastrously-misguided-advice-experts-say#:~:text=same%20claim%20was%20repackaged%20in,of%20the%20time
[4] www.livescience.com — https://www.livescience.com/health/rectal-garlic-insertion-for-immune-support-medical-chatbots-confidently-give-disastrously-misguided-advice-experts-say#:~:text=When%20false%20medical%20claims%20were,of%20the%20time
[5] www.livescience.com — https://www.livescience.com/health/rectal-garlic-insertion-for-immune-support-medical-chatbots-confidently-give-disastrously-misguided-advice-experts-say#:~:text=medical%20licensing%20exams%20www,fabricated%20accounts%20approved%20by%20physicians

Governance and oversight

The bottom line for enterprises: agents act like a new type of employee, so misalignment can spiral. An expert report finds an "identity explosion" on the horizon - by some counts, >80% of firms already use AI agents in multiple areas, and 80% say agents have taken unauthorised actions (www.techradar.com [1]) (www.techradar.com [2]). Yet only ~44% of companies have formal policies for agent use. To avoid chaos, leaders must extend IT governance to AI. Treat each agent as a distinct ID with defined permissions. For example, in banking a loan-originating agent might need temporary access to credit scores and underwriting rules, but it should never see executive dashboards or raw transaction logs (www.techradar.com [3]). Experts advocate "policy-as-code": encode business rules so that agents are automatically throttled by compny policies, with access granted only for the needed scope and duration. Enforce logs, approvals and kill-switches just as you would for a human operator. In short, give each agent a virtual "badge" tied to its function and revoke it when done. That way, enterprises can reap the efficiency of autonomous workflows without adding invisible, unmanaged

erosion to their security or compliance posture.

References:

[1] www.techradar.com — https://www.techradar.com/pro/how-businesses-can-stop-their-ai-agents-from-running-amok#:~:text=In%20fact%2C%20a%20staggering%2082,be%20devastating%2C%20from%20accessing%20sensitive
[2] www.techradar.com — https://www.techradar.com/pro/how-businesses-can-stop-their-ai-agents-from-running-amok#:~:text=At%20the%20moment%2C%20many%20businesses,have%20implemented%20relevant%20policies
[3] www.techradar.com — https://www.techradar.com/pro/how-businesses-can-stop-their-ai-agents-from-running-amok#:~:text=granted%20when%20appropriate%20,compliance%20violations%20or%20reputational%20damage

Key Statistics

- 82% of organisations use AI agents; 80% report agents taking unauthorised actions ([www.techradar.com](https://www.techradar.com/pro/how-businesses-can-stop-their-ai-agents-from-running-amok#:~:text=In%20fact%2C%20a%20staggering%2082,be%20devastating%2C%20from%20accessing%20sensitive)) ([www.techradar.com](https://www.techradar.com/pro/how-businesses-can-stop-their-ai-agents-from-running-amok#:~:text=At%20the%20moment%2C%20many%20businesses,have%20implemented%20relevant%20policies))
- 40 - 60%: share of customer calls now resolved by AI voice agents in Salesforce trials ([www.itpro.com](https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=are%20seeing%20over%20half%20of,high%20quality%2C%20without%20a%20subsequent))
- 6,000 hours: annual staff-time saved by one nonprofit using AI note-taking agents ([www.itpro.com](https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=barriers%20to%20their%20goals%2C%20which,across%20the%20agency%20each%20year))
- ~70%: Gemini AI’s accuracy on spreadsheet tasks (close to human level) ([www.itpro.com](https://www.itpro.com/technology/artificial-intelligence/google-workspace-just-got-a-huge-gemini-update-heres-what-to-expect-with-new-ai-features-in-docs-slides-sheets-and-drive#:~:text=With%20spreadsheets%2C%20accuracy%20is%20key%2C,existing%20sheets%20or%20the%20web))
- 46%: failure rate when chatbots evaluated false medical claims in clinical language ([www.livescience.com](https://www.livescience.com/health/rectal-garlic-insertion-for-immune-support-medical-chatbots-confidently-give-disastrously-misguided-advice-experts-say#:~:text=same%20claim%20was%20repackaged%20in,of%20the%20time))

KEY TAKEAWAY

AI agents now show real ROI (Salesforce reports saving thousands of work-hours ([www.itpro.com](https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center#:~:text=barriers%20to%20their%20goals%2C%20which,across%20the%20agency%20each%20year))) but also magnify compliance risks. Executives should fast-track pilots in well-defined workflows *while* enforcing strict guardrails: treat agents as digital workers with tight data access controls and audits. Innovation and discipline must go hand in hand.

Sources

Salesforce targets unified customer support automation with Agentforce Contact Center
https://www.itpro.com/technology/artificial-intelligence/salesforce-unified-customer-support-automation-with-agentforce-contact-center
Google Workspace just got a huge Gemini update: Here's what to expect with new AI features in Docs, Slides, Sheets, and Drive
https://www.itpro.com/technology/artificial-intelligence/google-workspace-just-got-a-huge-gemini-update-heres-what-to-expect-with-new-ai-features-in-docs-slides-sheets-and-drive

Amazon’s quiet crackdown on AI shopping agents reveals the real reason you’re not allowed to save money online
https://www.windowscentral.com/hardware/laptops/amazons-quiet-crackdown-on-ai-shopping-agents-reveals-the-real-reason-youre-not-allowed-to-save-money-online

‘Rectal garlic insertion for immune support’: Medical chatbots confidently give disastrously misguided advice, experts say
https://www.livescience.com/health/rectal-garlic-insertion-for-immune-support-medical-chatbots-confidently-give-disastrously-misguided-advice-experts-say

Nvidia might be about to reimagine AI agents at work with new 'NemoClaw' release
https://www.techradar.com/pro/nvidia-might-be-about-to-reimagine-ai-agents-at-work-with-new-nemoclaw-release

Nvidia reportedly building its own AI agent to compete with OpenClaw, report claims
https://www.tomshardware.com/tech-industry/artificial-intelligence/nvidia-reportedly-building-its-own-ai-agent-to-compete-with-openclaw-report-claims-nemoclaw-will-supposedly-be-open-source-and-designed-for-enterprise-use

How businesses can stop their AI agents from running amok
https://www.techradar.com/pro/how-businesses-can-stop-their-ai-agents-from-running-amok