

Slack & Microsoft Unveil Next-Gen AI Agents as Anthropic Leak Exposes Risks

Executive Summary

In the last 48 hours, leading tech vendors rolled out major AI agent enhancements while a major security slip-up raised alarm. Slack expanded its AI Slackbot via the Model Context Protocol to interface with thousands of enterprise apps (transcribing meetings and auto-populating CRM tasks) (www.itpro.com [1]). Microsoft broadened its Copilot AI suite: 'Copilot Cowork' (Anthropic Claude-powered workflow planner) joined the Frontier preview, and the Copilot 'Researcher' agent now uses multiple LLMs (OpenAI + Claude) to improve research accuracy (www.windowscentral.com [2]) (www.windowscentral.com [3]). Meanwhile, a bug caused Anthropic to expose ~500K lines of Claude's coding agent source code, revealing unreleased features and infrastructure (www.axios.com [4]). These developments signal big productivity gains from agents, but also new governance priorities (OpenClaw's 150K-star surge highlights employee-driven AI adoption). Executives should move quickly on pilots for efficiency — Slackbot users already save ~90 min/day (www.itpro.com [5]) — but shore up security and oversight against emerging risks.

References:

- [1] www.itpro.com — <https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=With%20expanded%20access%20to%20the,meeting%20actions>
[2] www.windowscentral.com — <https://www.windowscentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20recently%20announced%20that%20,Cowork%20to%20initiate%20multiple%20tasks>
[3] www.windowscentral.com — <https://www.windowscentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20has%20made%20major%20improvements,standard%20for%20deep%20research%20quality>
[4] www.axios.com — <https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=Anthropic%20leaked%20500%2C000%20lines%20of,its%20own%20source%20code>
[5] www.itpro.com — <https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added>

Slack's AI Slackbot Gets Smarter

Slack has significantly enhanced its AI-powered Slackbot by tapping into Salesforce's Model Context Protocol (MCP) (www.itpro.com [1]). In practice, this means Slackbot can "route user queries directly to Agentforce, as well as over 6,000 apps" within a company's tech stack (www.itpro.com [2]). Crucially, Slackbot can now sit in on virtual meetings (Zoom, Teams, etc.), transcribe what's said, and suggest follow-up actions – for example, creating calendar events or populating CRM fields after a call (www.itpro.com [3]). This effectively turns Slack into an agentic OS for work: Slackbot knows the full context of projects, channels, and files and can autonomously handle routine tasks.

Salesforce reports that Slackbot usage has escalated rapidly – on track to reach 1 million weekly users (the fastest feature adoption in Salesforce's 27-year history) (www.itpro.com [4]). Early adopters say Slackbot is saving them substantial time: some teams report roughly 90 minutes of work saved per day (www.itpro.com [5]). For example, Slackbot can take a meeting transcript and automatically

assign action items to team members or update enterprise systems, cutting hours of manual follow-up. The bottom line for leadership: these agent features make Slack more than just chat. Organizations should encourage power users to leverage Slackbot for routine workflows (report generation, data lookups, meeting summaries, etc.) to drive efficiency, while ensuring proper oversight of any automated actions.

Finally, Slack now offers "reusable skills" for Slackbot – essentially pre-built agent scripts – so organizations can roll out standardized bots across teams. Enterprise leaders should inventory common tasks (e.g. weekly reporting, approvals) that Slackbot can learn. Policies may need updates: with Slackbot able to access channels and files, firms should review access controls and ensure compliance monitoring extends to AI-driven actions.

References:

- [1] [www.itpro.com — https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=With%20expanded%20access%20to%20the,meeting%20actions](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=With%20expanded%20access%20to%20the,meeting%20actions)
- [2] [www.itpro.com — https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=With%20expanded%20access%20to%20the,meeting%20actions](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=With%20expanded%20access%20to%20the,meeting%20actions)
- [3] [www.itpro.com — https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slackbot%20can%20sit%20in%20on,meeting%20actions](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slackbot%20can%20sit%20in%20on,meeting%20actions)
- [4] [www.itpro.com — https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added)
- [5] [www.itpro.com — https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added)

Microsoft's Copilot Cowork and Multi-Model Researcher

Microsoft also made a major move: it has rolled out an AI agent called Copilot Cowork (powered by Anthropic's Claude) to customers in its Frontier program (www.windowscentral.com [1]). Copilot Cowork is designed to handle multi-step work delegation: a user describes a goal ("Finalize the Q2 budget, design charts, and email it to finance"), and the agent creates a plan, executes steps across different apps (Excel, Teams, email), and provides progress updates. Under the hood, Copilot Cowork uses Claude's workflow orchestration skills and Microsoft's connectors. This marks a shift from single-turn chatbots to multi-action agents. Business customers can now opt in to try this experimental Copilot in M365 (slots open for "First Frontier Suite" subscribers). While currently a preview, it illustrates how Microsoft is automating higher-level tasks, with the promise that mundane coordination work (like scheduling, data consolidation, and routine analysis) can be offloaded from employees to AI.

At the same time, Microsoft upgraded its Copilot "Researcher" AI assistant, which integrates into Office apps. Researcher now accesses both OpenAI and Anthropic models synergistically (www.windowscentral.com [2]). In practice, one model drafts an answer while another critiques it – a technique Microsoft calls "model council" or "Critique." The result: Microsoft reports Researcher achieves about a 13.8% higher score on a deep-research quality benchmark (www.windowscentral.com [3]). For example, when answering a complex business question, Copilot can compare OpenAI and Anthropic responses side-by-side and refine the output. This multi-model approach reflects a broader strategy: rather than bet on one AI, enterprises can harness multiple specialized LLMs. For decision-makers, this implies that AI assistants will diversify – procurement and IT teams may need to manage multiple AI subscriptions. It also means slightly higher costs (Microsoft announced a new "E7" M365 tier ~\$99/user for full AI feature access) but potentially greater returns via improved accuracy and flexibility. Leaders should evaluate which tasks benefit most from either a single or combined-model approach and prepare to negotiate with multiple vendors accordingly.

Together, Slack and Microsoft developments signal that AI agents are becoming embedded in everyday work tools. CFOs and COOs should engage with IT to pilot these new Copilots: for example, test Copilot Cowork on routine projects and compare its output quality to human-driven processes. At the same time, firms should ask vendors for usage data (e.g. how many operations Copilot has automated) to measure ROI. Importantly, teams must be trained to "supervise" these agents: AI still makes mistakes, so establish review processes (e.g. validate Copilot's meeting actions). By doing so, organizations can capture major time savings while catching errors early.

References:

- [1] [www.windowcentral.com — https://www.windowcentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20recently%20announced%20that%20,Cowork%20to%20initiate%20multiple%20tasks](https://www.windowcentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20recently%20announced%20that%20,Cowork%20to%20initiate%20multiple%20tasks)
- [2] [www.windowcentral.com — https://www.windowcentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20has%20made%20major%20improvements,standard%20for%20deep%20research%20quality](https://www.windowcentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20has%20made%20major%20improvements,standard%20for%20deep%20research%20quality)
- [3] [www.windowcentral.com — https://www.windowcentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20has%20made%20major%20improvements,standard%20for%20deep%20research%20quality](https://www.windowcentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20has%20made%20major%20improvements,standard%20for%20deep%20research%20quality)

Open-Source Agents: The Rise of OpenClaw

Not all agentic AI in the enterprise comes from big vendors. A surprising grassroots movement is happening: OpenClaw, an open-source, self-hosted agent framework, has exploded in popularity (www.techradar.com [1]). OpenClaw (formerly ClawdBot) lets anyone create an autonomous AI agent that lives on their own hardware and connects to consumer messaging platforms (WhatsApp, Telegram, etc.). In early 2026 it surged to roughly 150,000 GitHub stars within weeks, even causing a run on Mac Mini laptops in Asia (www.techradar.com [2]). The project's appeal is that it keeps user data on-premises: OpenClaw can perform tasks like email writing, web browsing, and file management, all while data never leaves the local machine (www.techradar.com [3]).

For organizations, OpenClaw's popularity is a double-edged sword. On one hand, it signals employee demand for personalized AI helpers. Developers or power users might deploy their own agents to automate tasks (e.g. an agent that automatically checks project progress from internal tools). On the other hand, these DIY agents bypass corporate controls. Executives must recognize that not all agentic innovation will originate from IT-led projects: tech-savvy staff may spin up such tools quietly. The implication is that companies should either embrace this trend or regulate it. Possible steps: provide an approved self-hosted agent platform (on secured hardware) for internal teams, or create rules requiring disclosure of personal AI bots. Security teams should audit these agents too – because if OpenClaw can execute shell commands, a malicious or buggy agent could be dangerous. In sum, OpenClaw shows that AI agents are as much an entrepreneurial force as they are vendor products. Firms should monitor open-source and internal AI tool use, ensuring they align with enterprise architecture and security policies, or even sponsor own internal agent projects to stay ahead.

References:

- [1] [www.techradar.com — https://www.techradar.com/pro/best-hardware-options-for-deploying-openclaw#:~:text=settlng%20on%20its%20current%20name,data%20staying%20within%20your%20purview](https://www.techradar.com/pro/best-hardware-options-for-deploying-openclaw#:~:text=settlng%20on%20its%20current%20name,data%20staying%20within%20your%20purview)
- [2] [www.techradar.com — https://www.techradar.com/pro/best-hardware-options-for-deploying-openclaw#:~:text=settlng%20on%20its%20current%20name,data%20staying%20within%20your%20purview](https://www.techradar.com/pro/best-hardware-options-for-deploying-openclaw#:~:text=settlng%20on%20its%20current%20name,data%20staying%20within%20your%20purview)
- [3] [www.techradar.com — https://www.techradar.com/pro/best-hardware-options-for-deploying-openclaw#:~:text=settlng%20on%20its%20current%20name,data%20staying%20within%20your%20purview](https://www.techradar.com/pro/best-hardware-options-for-deploying-openclaw#:~:text=settlng%20on%20its%20current%20name,data%20staying%20within%20your%20purview)

Governance and Risk: Lessons from a Claude Code Leak

The week's other big headline was a cautionary one. Anthropic, the AI lab behind Claude, mistakenly pushed a file containing about 500,000 lines of "Claude Code" – the internal codebase for its AI coding assistant ([www.axios.com \[1\]](https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=Anthropic%20leaked%20500%2C000%20lines%20of,its%20own%20source%20code)) – to a public software repository. This dump included architecture details, dozens of feature flags, and baked-in capabilities that had not yet been released ([www.axios.com \[2\]](https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=What%20they%27re%20saying%3A%20,to%20an%20Anthropic%20spokesperson%2C%20including)). In effect, every competitor now has a detailed look at how Claude's agentic system is built. While no customer data was exposed, the leak is akin to accidentally open-sourcing your product roadmap: it erodes competitive advantage and raises trust issues.

For enterprise leaders, this incident underscores why AI adoption demands robust governance. Unlike traditional software, even AI companies can make surprising mistakes with intellectual property. The takeaway: due diligence on AI vendors is now more than licensing and cost. Ask potential AI partners how they secure their code and models, and plan for the possibility of shared vulnerabilities. Internally, it means setting up oversight on how agentic tools are tested and monitored. The Anthropic leak is an example of the "responsibility gap" experts warn about – organizations must not blindly trailblaze with AI without processes to catch slip-ups. In practice, this could mean quarterly audits of AI systems, predefined incident-response plans, and legal safeguards around AI output and IP. By pairing the productivity gains of agents (Slack reports 90 min/day saved ([www.itpro.com \[3\]](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added))) with a higher governance bar, companies can push forward confidently rather than reactively scrambling after a breach.

References:

- [1] [www.axios.com](https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=Anthropic%20leaked%20500%2C000%20lines%20of,its%20own%20source%20code) — <https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=Anthropic%20leaked%20500%2C000%20lines%20of,its%20own%20source%20code>
- [2] [www.axios.com](https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=What%20they%27re%20saying%3A%20,to%20an%20Anthropic%20spokesperson%2C%20including) — <https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=What%20they%27re%20saying%3A%20,to%20an%20Anthropic%20spokesperson%2C%20including>
- [3] [www.itpro.com](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added) — <https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added>

Key Statistics

- 1,000,000 users: Slack reports ~1M weekly Slackbot users (fastest feature adoption in Salesforce history) ([\[www.itpro.com\]\(https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added\)](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added))).
- 90 minutes: Average daily work saved per person using Slackbot in early deployments ([\[www.itpro.com\]\(https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added\)](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added))).
- 13.8% improvement: Copilot Researcher (multi-model) scored 13.8% higher on the deep-research DRACO benchmark ([\[www.windowscentral.com\]\(https://www.windowscentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20has%20made%20major%20improvements,standard%20for%20deep%20research%20quality\)](https://www.windowscentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai#:~:text=Microsoft%20has%20made%20major%20improvements,standard%20for%20deep%20research%20quality))).
- 500,000 lines: Source-code leak in Anthropic's Claude Code exposed ~500K lines of internal agent code ([\[www.axios.com\]\(https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=Anthropic%20leaked%20500%2C000%20lines%20of,its%20own%20source%20code\)](https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai#:~:text=Anthropic%20leaked%20500%2C000%20lines%20of,its%20own%20source%20code))).
- 150,000 stars: GitHub popularity of OpenClaw (self-hosted agent) drove ~150K stars within weeks ([\[www.techradar.com\]\(https://www.techradar.com/pro/best-hardware-options-for-deploying-op-enclaw#:~:text=settlng%20on%20its%20current%20name,data%20staying%20within%20your%20p](https://www.techradar.com/pro/best-hardware-options-for-deploying-op-enclaw#:~:text=settlng%20on%20its%20current%20name,data%20staying%20within%20your%20p)urview)).

KEY TAKEAWAY

CEOs: AI agents are moving fast from hype into daily workflows (Slack reports ~90 min/day saved ([[www.itpro.com](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added)])(<https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access#:~:text=Slack%20reports%20high%20user%20uptake,per%20day%2C%20the%20firm%20added>))). Pilot them for routine tasks now, but couple each deployment with strict oversight. Vet AI vendors carefully (the Claude code leak is a warning), tighten data/IP policies, and clarify who manages AI-driven outcomes. In short: accelerate agents' gains, but reinforce governance to manage new risks.

Sources

[Slackbot just got a big update with MCP and desktop access](https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access)

<https://www.itpro.com/software/slackbot-just-got-a-big-update-with-mcp-and-desktop-access>

[This is Microsoft's new "Copilot Cowork": An experiment with Anthropic's Claude AI models that plans and delegates your work](https://www.windowscentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai)

<https://www.windowscentral.com/artificial-intelligence/microsoft-copilot/this-is-microsoft-new-copilot-cowork-ai>

[Anthropic leaked 500,000 lines of its own source code](https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai)

<https://www.axios.com/2026/03/31/anthropic-leaked-source-code-ai>

[Best hardware options for deploying OpenClaw](https://www.techradar.com/pro/best-hardware-options-for-deploying-openclaw)

<https://www.techradar.com/pro/best-hardware-options-for-deploying-openclaw>

