Perplexity announces “Computer,” an AI agent that assigns work to other AI agents

Perplexity has introduced “Computer,” a new tool that allows users to assign tasks and see them carried out by a system that coordinates multiple agents running various models.

The company claims that Computer, currently available to Perplexity Max subscribers, is “a system that creates and executes entire workflows” and “capable of running for hours or even months.”

The idea is that the user describes a specific outcome—something like “plan and execute a local digital marketing campaign for my restaurant” or “build me an Android app that helps me do a specific kind of research for my job.” Computer then ideates subtasks and assigns them to multiple agents as needed, running the models Perplexity deems best for those tasks.

The core reasoning engine currently runs Anthropic’s Claude Opus 4.6, while Gemini is used for deep research, Nano Banana for image generation, Veo 3.1 for video production, Grok for lightweight tasks where speed is a consideration, and ChatGPT 5.2 for “long-context recall and wide search.”

This kind of best-model-for-the-task approach differs from some competing products like Claude Cowork, which only uses Anthropic’s models.

All this happens in the cloud, with prebuilt integrations. “Every task runs in an isolated compute environment with access to a real filesystem, a real browser, and real tool integrations,” Perplexity says.

The idea is partly that this workflow was what some power users were already doing, and this aims to make that possible for a wider range of people who don’t want to deal with all that setup. People were already using multiple models and tailoring them to specific tasks based on perceived capabilities, while, for example, using MCP (Model Context Protocol) to give those models access to data and applications on their local machines. Perplexity Computer takes a different approach, but the goal is the same: have AI agents running tailor-picked models to perform tasks involving your own files, services, and applications.

Then there is OpenClaw, which you could perceive as the immediate predecessor to this concept.

The story so far

If you haven’t been following the wild OpenClaw craze, here’s the quick summary: originally titled ClawdBot, then Moltbot, OpenClaw was an agentic AI tool that leveraged large language models to independently operate as a sort of background or ambient process on your local machine, performing a wide range of tasks from sorting through your email history to building websites to, well, basically whatever you could imagine.

https://arstechnica.com/ai/2026/02/perplexity-announces-computer-an-ai-agent-that-assigns-work-to-other-ai-agents/