Anthropic debuts new model with hopes to corner the market beyond coding

Anthropic's "smartest model" is getting a major boost, the company said in a blog post announcing Claude Opus 4.6.

It called the new model a "direct upgrade" from its predecessor in a release, noting that it can better take on complex, multi-step tasks and get "much closer to production-ready quality on the first try than what we've seen with any model – documents, spreadsheets, and presentations will need less back-and-forth on iterations." It's available starting today with the same pricing as its predecessor, and according to the company, its particular strengths are in agentic coding, tool use, search, and financial analysis.

But most of all, it seems that with this release, Anthropic wants to expand Claude's current hype beyond just coding and corner the market on other types of knowledge work. With Opus 4.6, it invested in making the model better at creating presentations in PowerPoint and documents in Excel. The blog post included a plug for Cowork, Anthropic's recent release that's a non-tech-worker-friendly version of Claude Code, in hopes that users in non-technical industries will explore the use cases for research, marketing, and more.

On the coding front, Anthropic said in a release that Opus 4.6 was built to improve developers' experience with Claude Code even further, since it specializes in long-horizon tasks and can "take a development project that would normally take days and finish it in hours, handling everything from architecture to deployment."

The company also announced a feature currently in research preview called “agent teams,” allowing the new model to work within Claude Code “the way a real engineering team does,” meaning it’s possible to split one project’s work across agents that each own a part of the project and coordinate with each other.

Dianne Na Penn, Anthropic’s head of research product management, told *The Verge* that the company focused on bettering the “multi-agent” experience for developers with this launch, investing in output quality and speed, as well as getting the model to be better at other types of knowledge work besides just coding — i.e., Excel, PowerPoint, and search functions.

“This is the first version of an Opus model where we have a one-million context window offered in beta,” Penn said. “We just had such positive feedback about Opus 4.5 that one of the key features people wanted was a longer context window so they could work with Claude across more documents.”

Anthropic said in the blog post that it ran the “most comprehensive” set of safety tests for Opus 4.6 of any of its models to date. New evaluations included ones for user well-being, more complex tests on whether Opus 4.6 could refuse “potentially dangerous requests,” and updated tests for how well the model could secretly perform harmful actions. The model also displays heightened cybersecurity abilities, per the company, so it included six new cybersecurity probes to track potential misuse.

https://www.theverge.com/ai-artificial-intelligence/874440/anthropic-opus-4-6-new-model-claude