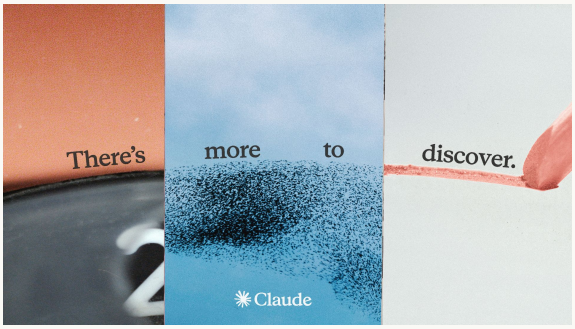
Claude BY ANTHROP\C

Sonnet 5.5: Smarter, faster, more efficient

Introducing Claude Sonnet 5.5, the second model in the Claude 5.5 family. It's a clear upgrade over Claude Sonnet 5, runs 30%+ faster, and costs up to 30% less for most work. Built for coding and knowledge work, Claude Sonnet 5.5 is a faster, lower-cost complement to Claude Opus 5.5 for well-scoped tasks.



Agents

Claude Sonnet 5.5 can take on well-defined agent tasks. It handles investigation, review, and drafting at a cost that makes it practical to do repeatedly.

Coding

Handles well-scoped software tasks like building features, reviewing code, and fixing bugs, and explains its changes clearly. Built for focused, everyday development.

Enterprise workflows

Brings improvements across knowledge work and office tasks, from analysis to documents and slide decks. Its writing is clearer and more direct, so drafts need less cleanup.

Everyday chat

A fast partner for everyday questions and brainstorms. Use it to power the assistants your teams and customers rely on, with quick, direct answers.

Financial analysis

Handles high-volume financial work, from summarizing filings and earnings to drafting everyday analysis, at a cost that works at scale.

Design and visual work

Turns prompts into slides, UI, diagrams, and user flows that are clearly structured and ready to refine.

When to use what model

Model	Use cases
Claude Fable 5.1	Frontier problems: long-running agents, net-new use cases, hardest coding & reasoning, full autonomy
Claude Opus 5.5	Your daily driver for complex work. Reasoning-intensive enterprise work (legal, financial analysis, complex knowledge work), complex and production coding
Claude Sonnet 5.5	A more cost-effective alternative to Opus for well-scoped coding and knowledge work: building features, fixing bugs, short polished documents, decks, and spreadsheets
Claude Haiku 4.5	High-frequency, latency-sensitive tasks; sub-agents in orchestration

Early users reported Claude Sonnet 5.5 does better work than Sonnet 5 while using fewer tokens and finishing faster. Across coding, customer support, finance, and knowledge work, teams saw it make fewer wrong decisions, recheck figures against source documents, and complete multi-step tasks that held up when the results were checked.

Why upgrade from Sonnet 5

Smarter

Anthropic's best Sonnet yet on coding and knowledge work, and a clear upgrade for anyone using Sonnet 5 today.

Faster

Sonnet 5.5 is 30%+ faster than Sonnet 5, making it Anthropic's fastest Sonnet model to date.

More efficient

Sonnet 5.5 typically needs far fewer tokens for the same work. In Anthropic's testing, it costs up to 30% less to run.

	Sonnet 5.5	Sonnet 5	Opus 5.5	GPT-6 Sol
Agentic coding Terminal-Bench 4.0	**70.6%**	10.3%	66.4%	—
Agentic coding FrontierCode 1.1 (Main)	46.2% max 52.1% xhigh	42.4%	**54.4%**	49.3%
Agentic coding CursorBench 4.0	55.5%	34.1%	**57.8%**	—
Knowledge work GDPval-AA v2.1	1844	1449	**1846**	1487
Knowledge work AA-Briefcase v1.1	1811	1359	**1822**	1483
Multidisciplinary reasoning Humanity's Last Exam (with tools)	64.5%	54.9%	**67.7%**	—
Computer use OSWorld 2.1 (partial)	80.1%	57.0%	**81.8%**	—
Visual chart recognition Chartography (no tools)	61.6%	15.6%	**64.4%**	53.6%

Stronger safeguards

Anthropic is continually improving its safety practices to reliably catch and prevent risks Claude models might pose, particularly as they become much more capable. Claude Sonnet 5.5 shows a substantial improvement in scores across Anthropic's alignment evaluations relative to Claude Sonnet 5, and comes with new, more robust safeguards that are proportionate to its leap forward in capabilities.