

# Setting the standard for enterprise AI

**Claude Opus 4.8** is our most intelligent Opus model and the best generally available model for coding and agents, with deeper reasoning for enterprise workflows. It excels in areas that require precision across long horizons, including financial analysis, cybersecurity, and computer use.



## What's new

**Fast mode.** Now runs on Opus 4.8 at \$10/\$50 per MTok, a 66.7% price drop from fast mode for Opus 4.7. Get the same Opus 4.8 intelligence at ~2.5x output speed for 2x the standard rate. Available in Claude Code (`/fast on`) and Claude Developer Platform (research preview).

**Mid-conversation system messages** (research preview). Developers can update Claude's instructions mid-task without breaking the prompt cache or routing the update through a user turn. Use this in your harness to update permissions, token budgets, or environment context as your agent runs.

## Built for how organizations actually work

Opus 4.8 changes what teams can hand off to Claude. It follows instructions, stays in scope, and produces outputs teams can ship without re-checking. Lower output variance and stronger instruction-following mean fewer guardrails, less post-processing, and fewer review cycles.



### Long-running agents

Works independently for longer, unblocking itself across multi-stage projects and complex dependency chains with minimal oversight.



### Coding

Keeps track of dependencies, gets unstuck on its own, and persists over hours on long, multi-stage tasks with stronger memory and longer context.



### Enterprise workflows

Reasons across long documents, checks its work, and manages multi-stage projects end-to-end with professional polish on spreadsheets, slides, and docs.



### Computer use

Pairs improved vision with deep reasoning for multi-step tasks that span multiple applications and require planning and judgment.



### Financial analysis

Brings deeper reasoning and precision to investment research, earnings analysis, carrying context across an entire reporting cycle, and more.



### Cybersecurity

Brings deeper reasoning to security workflows like threat intelligence synthesis, vulnerability finding, alert triage, and incident response.



Claude Opus 4.8 has noticeably better judgment. In Claude Code, it asks the right questions, catches its own mistakes, pushes back when a plan isn't sound, and builds up confidence around complex, multi-service explorations before making big changes. It's a great model to build with.

– Tom Prichard, Staff Engineer

## Opus 4.8 in Claude Code

In Claude Code, Opus 4.8 further refines agentic software development. It makes calls like an experienced engineer without needing constant check-ins, stays on track across long-running sessions, and follows work through entire repos, allowing you to hand off a feature, a migration, or a bug sweep so teams can focus on what's next.

### Dynamic workflows (research preview)

Alongside Opus 4.8, we're releasing dynamic workflows, a new feature that lets Claude take on larger tasks end to end, planning the work, running hundreds of parallel subagents, and verifying what comes back. Available in Claude Code for Enterprise, Teams, and Max plans.

### Frontier intelligence for deeper knowledge work

	Opus 4.8	Opus 4.7	GPT-5.5	Gemini 3.1 Pro
<b>Agentic coding</b> SWE-bench Pro	69.2%	64.3%	58.6%	54.2%
<b>Agentic terminal coding</b> Terminal Bench 2.1	74.6%	66.1%	78.2%	70.3%
<b>Multi-disciplinary reasoning</b> Humanity's Last Exam	49.8% without tools 57.9% with tools	46.9% without tools 54.7% with tools	41.4% without tools 52.2% with tools	44.4% without tools 51.4% with tools
<b>Agentic computer use</b> OSWorld-Verified	83.4%	82.8%	78.7%	76.2%
<b>Knowledge work</b> GDPval-AA	1890	1753	1769	1314
<b>Agentic financial analysis</b> Finance Agent v1.1	53.9%	51.5%	51.8%	43.0%

### Effort in Claude.ai and Claude Cowork

A new control alongside the model selector lets you choose how much effort Claude puts into a response: on lower effort Claude will respond faster and burn down rate-limits slower, while on higher effort gives Claude will think more deeply for harder problems. Available on all plans.

