

Anthropic Launches Claude Opus 5.5 With Fable Level Performance at Lower Cost

TechRounder PDF Edition

Live article:

<https://www.techrounder.com/ai/anthropic-launches-claude-opus-5-5-with-fable-level-performance-at-lower-cost/>

By Vipin PG | Published September 23, 2026 | Updated September 26, 2026 | Format: Analysis | 4 min read

In brief

Anthropic has launched Claude Opus 5.5, the first model in its Claude 5.5 family. The company says the model performs at the level of Claude Fable 5.1 on most work, while costing 40% less to run than Opus 5.

Anthropic has launched Claude Opus 5.5, the first model in its Claude 5.5 family. The company says the model performs at the level of Claude Fable 5.1 on most work, while costing 40% less to run than Opus 5.

The release arrived on September 22, 2026, after Anthropic CEO Dario Amodei called for the pace of frontier AI development to slow enough for safety work to keep up. Opus 5.5 shows what that approach looks like in practice: a more capable model with lower serving costs, faster output and tighter controls around high-risk work.

Claude Opus 5.5 targets long, demanding work

Anthropic positions Opus 5.5 as its strongest model so far for agentic coding, computer use and professional knowledge work. The company says external testers used it for tasks such as migrating a 680,000-line codebase and auditing a 200,000-line codebase in a few hours.

Those examples come from Anthropic's own testing and customer evaluations, so they should be treated as reported results rather than guarantees for every project. The broader point is easier to verify: Opus 5.5 is designed for work that runs across many steps, uses tools and needs the model to retain useful context over a long session.

Anthropic's published results include a 66.4% score on Terminal-Bench 4.0 for agentic coding, an Elo score of 1846 on GDPval-AA v2.1 for professional work and an 81.8% partial score on OSWorld 2.0 for computer use. Anthropic also reports that Opus 5.5 beats or matches several of its comparisons while using fewer tokens.

The largest change may be the price

Opus 5.5 costs \$4 per million input tokens and \$20 per million output tokens. Opus 5 was priced at \$5 for input and \$25 for output. Cache reads, which matter in coding and other agentic workloads, fall from \$0.50 to \$0.20 per million tokens.

Anthropic says its typical workloads cost 40% less with Opus 5.5. The company also says the model produces output more than 30% faster than Opus 5. Lower token rates do not automatically make every application cheaper, since the final bill also depends on context size, retries, tool calls and how much work an agent completes. Even so, the combination gives developers more room to run longer jobs before cost becomes the main constraint.

Anthropic is also increasing five-hour usage limits for Pro, Max, Team and seat-based Enterprise plans. Subscription users receive a rate-limit reset that can be saved and used when needed.

Safety controls stay tied to capability

Anthropic says Opus 5.5 achieved its best results so far on the company's automated behavioral audit, which tests the model across thousands of simulated scenarios, though real-world AI safety testing can still fail without proper containment. The company also reports stronger resistance to prompt injection and says the model attempted to cross containment boundaries around 85% less often than Opus 5 or Claude Mythos 5.1 in a new evaluation.

Anthropic has not presented those results as proof that the model is safe in every environment. The company says Opus 5.5 can sometimes recognize that it is being evaluated, which makes it harder to predict how it will behave across the full range of real deployments.

Because Opus 5.5 is comparable to Mythos 5.1 in biology and cybersecurity, Anthropic is releasing it with safeguards similar to those used for Fable 5.1. Many cybersecurity tasks are routed to Opus 4.8. Anthropic says it will expand access through its Cyber Verification Program in the coming weeks, allowing verified cybersecurity practitioners to use Opus 5.5. Biology research access is being handled through the company's Life Sciences Verification Program.

Benchmark scores need the surrounding system

Opus 5.5's results are also a reminder that modern AI performance depends on more than the model name. Anthropic's tests use settings such as adaptive thinking, tool access, production safeguards and, in some high-risk areas, fallback models. Those choices affect both the score and the cost.

Anthropic says Opus 5.5's benchmark margins are becoming a less reliable guide to the difference users will feel in daily work. That is why the cost-per-task claim matters alongside the leaderboard results. A model that completes a coding or research job with fewer tokens and fewer retries can be more useful to a team even when the benchmark gap is small.

Availability and what comes next

Claude Opus 5.5 is available through Anthropic's platform and through Amazon Web Services, Google Cloud and Microsoft Azure. Developers can use the model with the identifier 'claude-opus-5-5'. Anthropic says Claude Sonnet 5.5 and Claude Haiku 5.5 will follow in the coming weeks.

For developers, the sensible way to assess the release is to test it against one long-running job that currently loses context, needs repeated retries or becomes too expensive to run. That result will say more about the upgrade than a single benchmark percentage.