

# Claude Fable 5

## Anthropic Opens the Door to a Mythos-Class Model

---

The most capable model Anthropic has ever shipped to the public — and what our early testing says about how to actually use it.

# One model, two faces

On June 9, Anthropic did what it had hesitated to do for two months: put a Mythos-class model in everyone's hands. Fable 5 and Mythos 5 are the same model underneath — the difference is the safeguards that make Fable safe to ship at scale.

Even the name carries the idea: Fable, from the Latin fabula — “that which is told” — a deliberate echo of the Greek mythos.

## CLAUDE FABLE 5

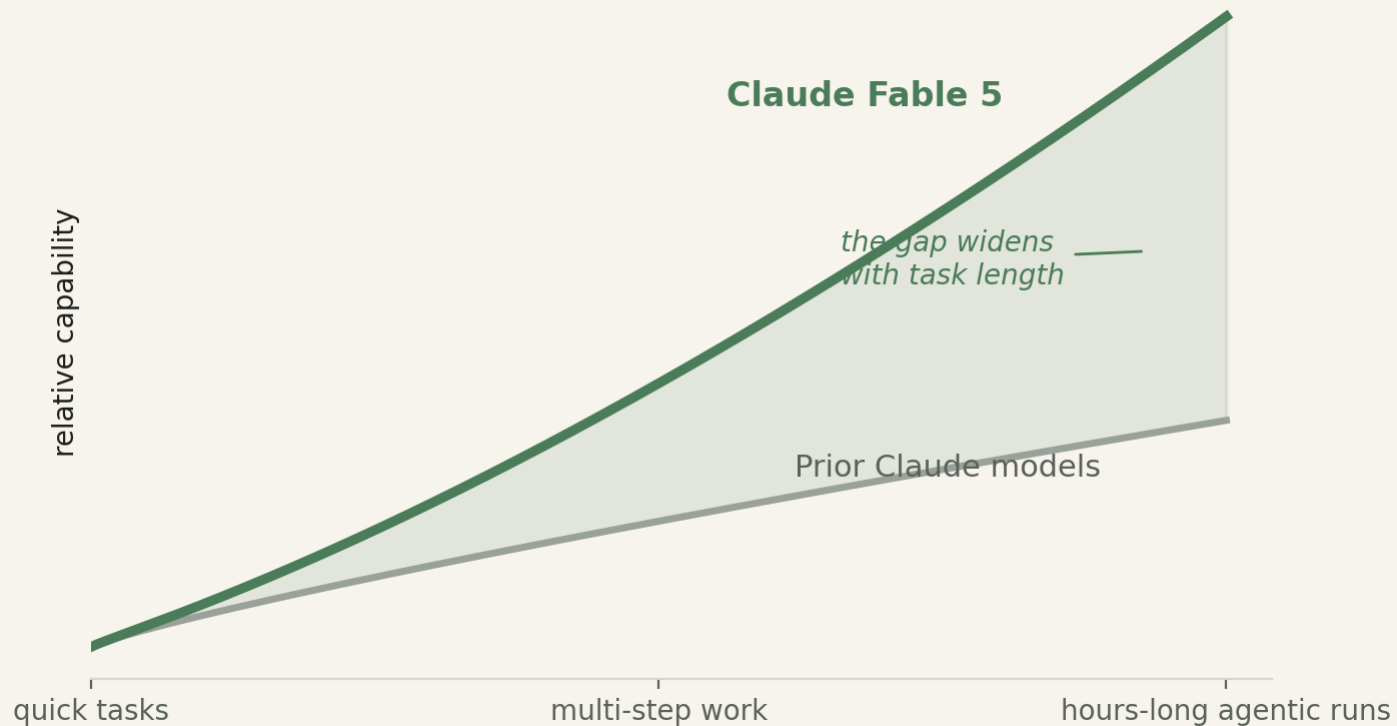
Generally available to everyone  
Safety classifiers on dual-use domains  
The cautious public storyteller

## CLAUDE MYTHOS 5

Gated to approved organizations  
No capability ceiling  
The unguarded specialist

# It doesn't win on one-liners. It wins on endurance.

State-of-the-art on nearly every tested benchmark. The line worth underlining: the longer and more complex the task, the larger Fable 5's lead grows.



# Launch day, loud and clear

“ **Andrej Karpathy** @karpathy

A step change worthy of a major version bump. Hand it far more ambitious tasks — the model “gets it” and runs with them. (Safeguards feel a touch trigger-happy at launch.)

“ **Boris Cherny** @bcherny · built Claude Code

The best model he's used “by a wide margin” — a real step up in code quality, tool use, and self-verification, with fewer prompts and steers needed.

“ **Afolabi Sokeye** @SokeyeA

“Oh we are so up.” Plus a practical scramble: the free window is short, so users rushed to test it before the terms change.

# The bottleneck isn't the model's stamina anymore. It's your scoping.

- Genuinely good at code — specifically at large-scale work. It runs for hours without losing the thread.
- The unlock is the verification loop: have the model check, test, and correct its own output.
- That's not a marginal bump. It's the difference between a draft and something you'd ship.
- Fable seems built to reward that pattern rather than fight it.

*Preliminary — we're still in early testing. But this is the pattern that keeps showing up.*

# Intelligence demand is infinite

As capable models get cheaper and sustain longer, harder work, the constraint shifts from “can it do this?” to “what's the next ambitious thing we'll ask it to do?” — and that list never ends.

Karpathy reached the same idea via Jevons paradox: as working software starts coming out “on a tap,” his appetite for it grew rather than shrank.

**Explainers. Visualizers. Dashboards. Bespoke single-use apps. 10x'd test suites. Whole research projects on demand.**



**demand for  
intelligence**

Read the companion essay:  
Intelligence Demand Is Infinite  
on [getcoai.com](https://getcoai.com)

# Developers and builders are the headline audience

Built for agentic, long-horizon coding workflows — and more token-efficient than past Claude models, which matters for cost as much as speed.

**Codebase-wide migration: ~50M lines of Ruby**

Traditional engineering team



**2+ months**

Claude Fable 5 **1 day**



# Knowledge work and vision

## KNOWLEDGE WORKERS

- Top scores on senior-level financial reasoning benchmarks
- Real gains in document-based reasoning
- Better at interpreting charts and tables — the unglamorous work that fills most professional days

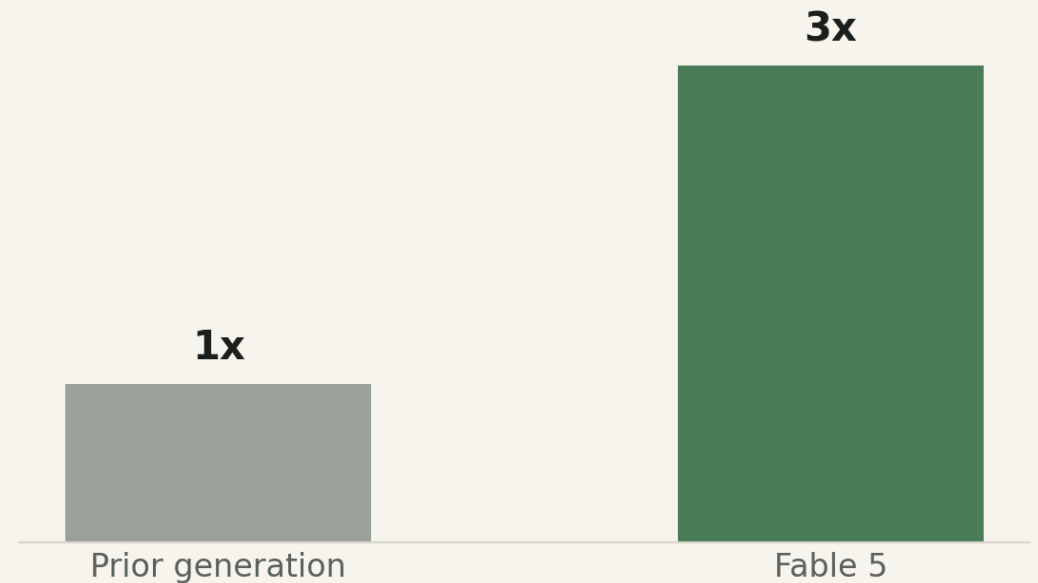
## IMAGES & SCREENS

- Pulls precise numbers out of dense scientific figures
- Rebuilds a web app's source code from screenshots alone
- Finished Pokémon FireRed using only raw game screenshots — no navigation scaffolding

# Long-context and memory-heavy tasks

Fable holds focus across millions of tokens and improves its own output using its notes. Given persistent memory in a deck-building game, its performance gains tripled versus the prior generation.

**Performance gains with persistent memory  
(deck-building game benchmark)**




# Who Fable 5 is not for (today)

- Queries touching cybersecurity, biology, chemistry, or model distillation are intercepted and quietly handed to Claude Opus 4.8.
- Safeguards are tuned conservatively — they sometimes catch harmless requests.
- Unrestricted capability stays with Mythos 5: gated to Project Glasswing cyber partners and, soon, approved biology researchers.

**For the vast majority of users, Fable behaves identically to the unrestricted Mythos 5.**





**API: claude-fable-5 · fully available at launch**

**\$10** per M input tokens

**\$50** per M output tokens

Less than half the cost of the earlier Mythos Preview — frontier capability priced to actually be paid for.

# The era of waiting behind a gated preview is over.

- Anthropic's strategy, stated plainly: build the most capable models possible, wrap them in safeguards strong enough to release responsibly, and get them into as many hands as possible.
- The guardrails are real, and some users will bump into them.
- But the trade is a frontier-class model — one that shines on the long, hard, multi-step work that used to take a whole team — available to anyone with an API key, at a price designed to be paid.

CO/AI

# All Signal. No Noise.

One concise email a day to make you smarter on AI —  
curated by Anthony Batt & Harry DeMott.

[Subscribe at getcoai.com](https://getcoai.com) →

## FOLLOW THE CONVERSATION

X / Twitter: [@getcoai](https://twitter.com/getcoai) · [@djabatt](https://twitter.com/djabatt) (Anthony Batt)

[Daily Briefings](#) · [Essays](#) · [CO/AI Podcast](#) · [Video Feed](#) — [getcoai.com](https://getcoai.com)