

RELEASED 9 JUNE 2026

Claude Fable 5 and Mythos: Everything You Need to Know

What Anthropic's most capable model means for your business, what it really costs to use, and what to do before the free window closes on 22 June 2026.

Turn Intent Into Impact.



One engine. Two doors.

On 9 June 2026 Anthropic released its most capable model ever as two products. Same underlying model, different doors. You get one of them, and you are missing nothing behind the other.

THE PUBLIC DOOR

Claude Fable 5

Generally available now in paid Claude.ai plans, Claude Code and the API.

- Safety classifiers watch for misuse in three areas
- Flagged sessions get quietly answered by Opus 4.8
- Under 5% of sessions affected, almost never business topics

"A Mythos-class model made safe for general use." Anthropic, 9 June 2026

THE RESTRICTED DOOR

Claude Mythos 5

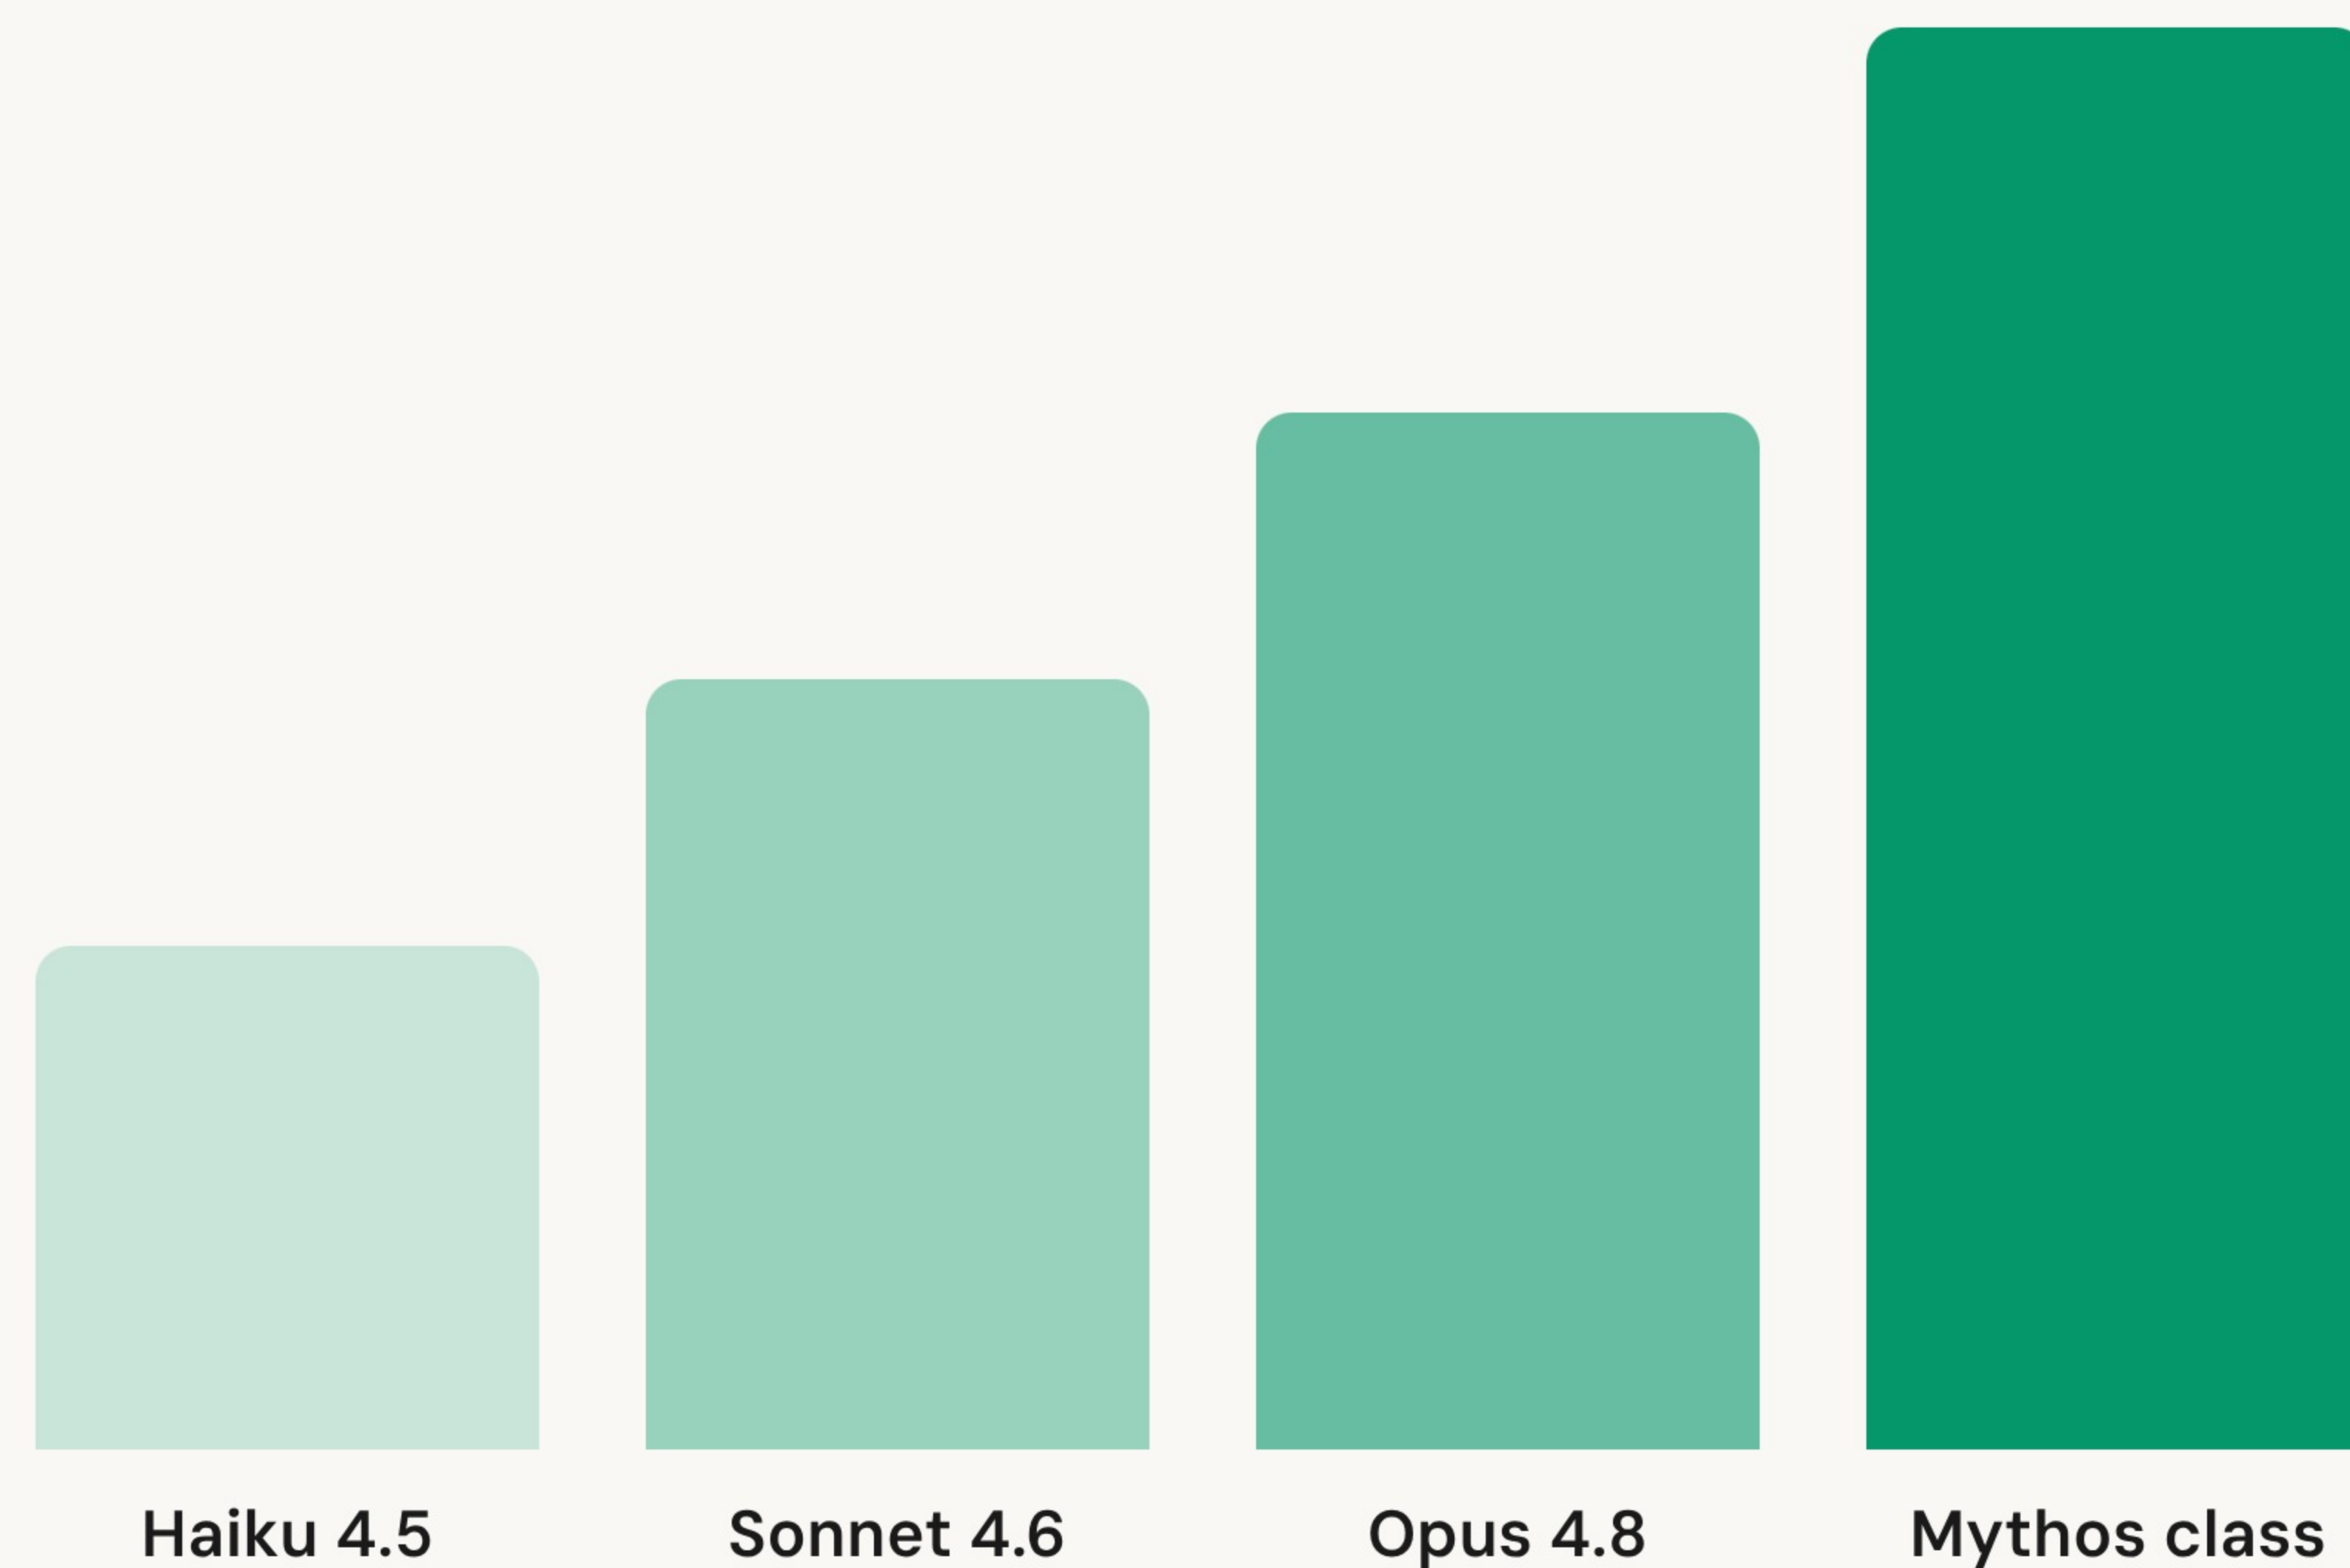
The identical model with those safeguards lifted. You cannot buy access.

- Roughly 150 vetted security organisations
- Project Glasswing, run in consultation with the US government
- Identical capability and identical price to Fable 5

If you run an SME, there is nothing behind this door you are missing.

Mythos-class: a new tier above Opus

Until now, Opus was the top of the Claude range. Mythos-class models sit above it. Fable 5 is the first one the public can use.



WHO CAN GET MYTHOS 5?

A vetting programme, not a product tier

- Existing Project Glasswing partners: critical-infrastructure and cyberdefence organisations
- Expansion planned for vetted biology researchers, in consultation with the US government
- No public application route, no price you can pay to get in

Why you should not care: it is the identical model at the identical price. Everything Mythos 5 can do for a business, Fable 5 can do too.

Double the price of Opus. More in practice.

API prices per million tokens, as published by Anthropic on 9 June 2026.

MODEL	INPUT	OUTPUT
Fable 5 / Mythos 5	\$10	\$50
Opus 4.8	\$5	\$25
Sonnet 4.6	\$3	\$15
Haiku 4.5	\$1	\$5

1M

token context window, as standard

128k

output tokens per request

+30%

more tokens for the same content under the new tokenizer, so the real cost is more than double Opus



Included free on every paid Claude plan until 22 June 2026. From 23 June it draws pay-as-you-go usage credits.

Built for long, autonomous work. Not better emails.

The gains show up when you hand it a whole job, not a single question. Three data points from the first 48 hours:

2 months, 1 day

Stripe compressed a two-month codebase migration into a single day, with serious engineering scaffolding around it.

Anthropic launch announcement

80.3%

SWE-Bench Pro coding score, against 58.6% for GPT-5.5. The Claude Code lead calls it the best coding model he has used, by a wide margin.

System card recap, launch coverage

2h vs 8h

First-hand: it wrote its own fuzzer to catch bugs its developer missed, and finished in 2 hours what took the previous model 8.

Hacker News launch thread



Honesty check: on Blueprint-Bench it scores 38.6% against 58.6% for humans. This is frontier, not superhuman.

Three tripwires. One quiet fallback.

Separate classifiers watch every Fable 5 session for three kinds of misuse. A flagged session gets answered by the older Opus 4.8 instead, silently. It happens in under 5% of sessions, and almost never on business topics.

≤ 5%

🛡️ Cybersecurity

🧪 Biology and chemistry

📄 Model copying (distillation)

FALSE POSITIVES ARE REAL

In week one: a medical MRI task flagged as bioterrorism, sales prospecting flagged as cybersecurity. If you work in IT security or life sciences, expect occasional friction. Everyone else will likely never notice.

ONE FLAG FOR REGULATED WORK

Fable 5 requires 30-day data retention and is not available under zero-data-retention agreements. If you handle client data under strict confidentiality commitments, check before adopting.

Six things to do before 22 June

01 Test your hardest task
Throw it at the messy backlog job: the spreadsheet consolidation, the contract-pack review. Long multi-step work is its edge.

02 Run a side by side
Same prompt to your current model and to Fable 5. No visible difference on your work means stay on the cheaper tier.

03 Set a spending cap
When the usage window empties, Claude can keep going on credits charged to your card. Cap it in account settings now.

04 Verify before you trust
Its own system card documents fabricated verification claims in testing. Spot-check numbers and totals like a new hire's work.

05 Put 22 June in the calendar
So nobody on the team unknowingly burns paid credits on 23 June for tasks the cheaper models handle fine.

06 Building on the API? Hold.
Anthropic keeps Opus 4.8 as the standard upgrade path. Reserve Fable 5 for tasks Opus demonstrably fails.

The cost trap is real.

Per finished task it can be cheaper than Opus. Per hour of heavy agentic use, it is brutal. First-hand reports from week one:

\$160/hr

sustained burn rate on heavy agentic coding, measured by one developer

9 min

to drain a \$100 daily allowance, reported by one early tester

\$1,000+

spent in a single day by one developer on a \$200 monthly plan



Set a hard spending cap before you experiment. When the usage window empties, Claude keeps going on credits charged to your card.

Four things you can safely ignore

Mythos envy

You cannot get it, and it is the identical model at the identical price. There is nothing on the other side of that door for an SME.

Benchmark theatre

Launch coverage quotes the higher of the two models' scores in one combined column. The public Fable 5 number can be lower.

The overnight-game demos

The viral builds ran 3 to 4+ hours of continuous compute. Impressive, not representative of a \$20 per month workflow.

Doom headlines

The 'AI is getting too dangerous' debate is about frontier-lab coordination, not about risk to your business using the product.

Every claim in this guide is traceable

Anthropic	Introducing Claude Fable 5 and Claude Mythos 5	anthropic.com/news/claude-fable-5-mythos-5 , 9 June 2026
Anthropic docs	Model introduction, pricing and safeguard mechanics	platform.claude.com , 9 June 2026
TechCrunch	Launch report including pricing and retention policy	techcrunch.com , 9 June 2026
Simon Willison	First-day hands-on testing in Claude Code and Claude.ai	simonwillison.net , 9 June 2026
The Neuron	Explainer including tokenizer, benchmarks and cost analysis	theneuron.ai , June 2026
Hacker News	Launch thread with first-hand practitioner reports	news.ycombinator.com , 9 to 11 June 2026
Developers Digest	The 22 June deadline and usage-credit mechanics	developersdigest.tech , June 2026
Anthropic	Project Glasswing partner programme	anthropic.com/glasswing , 7 April 2026

Test before you trust.

We help SMEs adopt AI with evidence, not hype. If you want a second pair of eyes on what Fable 5 could actually do for your business, talk to us before the free window closes.

Talk to us: bykovbrett.net/contact

More guides: bykovbrett.net/resources



Turn Intent Into Impact.