Claude Models Roadmap May 2026: Opus 4.7, Knowledge Cutoffs, the 1M Context Window, and What's Real About Claude 5

Edition of November 13, 2026 — facts verified November 13, 2026. Current edition always at https://tygartmedia.com/claude-models-roadmap-may-2026-opus-47-claude-5/.

Direct Answer (November 13, 2026): The May 2026 roadmap below has played out through multiple generations. Anthropic's current active lineup is **Claude Fable 5.2** (`claude-fable-5-2`, $8/$40), **Claude Opus 5.5** (`claude-opus-5-5`, $4/$20), **Claude Sonnet 5.5** (`claude-sonnet-5-5`, $2/$10), and **Claude Haiku 5.5** (`claude-haiku-5-5`, $0.10/$0.50 ≤100K) — all with a 1M context window and 128K output. Legacy June–September releases (Fable 5/5.1, Opus 4.8, Sonnet 5, Haiku 4.5) are archived in-place below for provenance. See current model tracker and Claude Fable guide.

Last refreshed: November 13, 2026 · Primary sources: Anthropic models overview, API pricing & Anthropic news.

The pace of new Claude releases in 2026 has been fast enough that the canonical question — "what's the latest Claude model and what's it actually good for?" — has a different answer almost every quarter. This article is the current map, dated and sourced, of what Anthropic has shipped in 2026, what's confirmed about each model's specs and knowledge cutoffs, and what's been claimed (but not officially confirmed by Anthropic) about what's coming next.

Two ground rules first, because the model-roadmap space is full of speculation:

- Specs and release dates marked as **verified** come from Anthropic's own documentation, news posts, or help center pages. We list the specific source.
- Anything marked as **reported** or **claimed** comes from third-party reporting (TechCrunch, secondary news sites, analyst commentary) that we could not independently confirm against an Anthropic-published source as of May 15, 2026.

If you're making product decisions on this information, treat verified facts as actionable and reported facts as directional.

The May 15, 2026 generally-available Claude models (historical snapshot)

May 2026 GA Claude models — historical snapshot framing.

Historical Snapshot (May 15, 2026) — Superseded by Nov 13, 2026 Lineup: Current active models on Anthropic's pricing schedule are **Fable 5.2 ($8/$40)**, **Opus 5.5 ($4/$20)**, **Sonnet 5.5 ($2/$10)**, and **Haiku 5.5 ($0.10/$0.50 ≤100K)** — all 1M context / 128K output (legacy Fable 5.1 $10/$50, Opus 4.8 $5/$25, Sonnet 5 $2/$10, Sonnet 4.6 $3/$15, Haiku 4.5 $1/$5). The bullets directly below preserve what was active on May 15, 2026:

Claude Opus 4.7 — `claude-opus-4-7`

- **Status:** Legacy — superseded (May 2026 model; replaced in the Opus tier by **Opus 5.5** and in the flagship slot by **Fable 5.2**).
- **Context window:** 1 million tokens at standard pricing (no long-context premium)
- **Max output:** 128,000 tokens
- **Knowledge cutoff:** January 2026 (per Anthropic Help Center, verified May 15, 2026)
- **Pricing:** $5/MTok input, $25/MTok output (base rates)
- **Notable changes from 4.6:** New tokenizer (uses up to ~35% more tokens for the same text), high-resolution image support up to 2576px / 3.75MP, new xhigh effort level, task budgets beta. Extended thinking budgets and sampling parameters (temperature, top_p, top_k) are removed.

Claude Opus 4.6 — Legacy (Feb 2026, $5/$25; succeeded by Opus 4.8, Opus 5, and now Opus 5.5 at $4/$20).

Claude Sonnet 4.6 — Legacy ($3/$15; succeeded by Sonnet 5 and now Sonnet 5.5 at $2/$10).

Claude Haiku 4.5 — Legacy (200K context, $1/$5; now superseded by Haiku 5.5 at $0.10/$0.50 ≤100K with 1M context).

Earlier models still active or in deprecation: Opus 4.5, Opus 4.1, Sonnet 4.5, and Haiku 3.5 (retired except on Bedrock and Vertex AI). Opus 4 and Sonnet 4 are listed as deprecated.

Knowledge cutoff dates that actually matter

Per Anthropic's Help Center article on training-data recency (verified May 15, 2026), the most recent generally-available models have January 2026 knowledge cutoffs. That means:

- Anything that happened after January 2026 is outside the model's training data
- For current events, recent product launches, recent legal or regulatory changes, or very recent technical documentation, the model needs to be given the information directly (in the prompt, via web search, or through tool use) — it can't be relied on to know it
- The model still has tools available (web search, code execution, file access) that can access post-cutoff information when explicitly invoked

The practical version: don't ask Claude what happened last week and expect it to know. Hand it the source material and ask it to analyze, summarize, or work with what you've given it.

The 1M token context window — what it actually unlocks

1M token context — what it unlocks.

Per Anthropic's official pricing documentation (verified May 15, 2026), Opus 4.7, Opus 4.6, and Sonnet 4.6 all include the full 1 million token context window at standard pricing. There's no long-context premium — a 900,000-token request is billed at the same per-token rate as a 9,000-token request.

That's an enormous practical change from earlier Claude generations. A 1M context window is roughly:

- ~750,000 words of English text
- Most full books or technical specifications in a single context
- ~8 hours of meeting transcripts at typical density
- An entire mid-sized codebase, including most or all source files

Prompt caching and batch processing discounts both apply across the full 1M context window on every current model (**Fable 5.2, Opus 5.5, Sonnet 5.5, and Haiku 5.5**).

What's reported about Claude 5 (and what we cannot independently verify)

What's reported about Claude 5 — verify carefully.

Multiple third-party sources reported in early 2026 that Anthropic CEO Dario Amodei confirmed a Q2 2026 launch window for Claude 5 in a TechCrunch interview published February 1, 2026. The same sources cited an internal-roadmap leak suggesting an April 28 target date.

What we can verify as of May 15, 2026:

- When this roadmap was written, Anthropic's documentation listed Opus 4.7 as the latest production model. That has since advanced: Opus 4.8 and the Fable 5 / Mythos 5 top tier shipped in June 2026, and Claude Sonnet 5 was released June 30, 2026 — so the "5" generation is now real by name, arriving as Sonnet 5 plus the Fable / Mythos tier rather than a single "Claude 5" release.
- The third-party reporting on Claude 5 specifications (500K context window, 20-25% benchmark improvements, ~90%+ on SWE-bench Verified) is widely repeated but, as far as we could verify, is not sourced to an Anthropic-published document.

Looking back from November 13, 2026, that 5.x wave fully arrived — advancing from Fable 5 / Opus 4.8 / Sonnet 5 in June to **Fable 5.2, Opus 5.5, Sonnet 5.5, and Haiku 5.5** today.

How the Claude 5 / 5.5 rollout played out

Timeline summary: Claude Sonnet 5 initially shipped June 30, 2026 (`claude-sonnet-5`, 64K output) alongside Fable 5 and Opus 4.8, and was later succeeded by **Sonnet 5.5** (`claude-sonnet-5-5`, $2/$10, 1M/128K), **Opus 5.5** (`claude-opus-5-5`, $4/$20), **Haiku 5.5** (`claude-haiku-5-5`, $0.10/$0.50 ≤100K), and **Fable 5.2** (`claude-fable-5-2`, $8/$40). The currently recommended production model IDs are `claude-fable-5-2`, `claude-opus-5-5`, `claude-sonnet-5-5`, and `claude-haiku-5-5`.

How to actually keep up with Claude releases

The signal-to-noise ratio in the model-release coverage space is not great. Two practical sources are reliable enough to bookmark:

- **Anthropic's news page** at anthropic.com/news — first-party launch announcements with full model details
- **Anthropic models overview** at platform.claude.com — concise, live model IDs, context limits, and retirement schedules

For breaking changes that affect production code, the Anthropic platform documentation publishes per-version "What's new" pages (Opus 4.7's, for example, lists every API breaking change at launch). Those are the canonical reference for migration work.

For everything else — analyst commentary, predictions, leak coverage — treat it as commentary, not as fact.

What this means for your work today

Based on Anthropic's verified lineup as of November 13, 2026:

- For frontier reasoning or long-horizon agents, use **Fable 5.2** (`claude-fable-5-2`, $8/$40, 1M context).
- For primary software engineering and daily default reasoning, use **Opus 5.5** (`claude-opus-5-5`, $4/$20, 1M context).
- For cost-efficient production workloads, use **Sonnet 5.5** (`claude-sonnet-5-5`, $2/$10, 1M context / 128K output).
- For high-throughput, low-cost tasks, use **Haiku 5.5** (`claude-haiku-5-5`, $0.10/$0.50 ≤100K, 1M context).
- For current events, SEC/IPO filings, or post-cutoff docs, pass primary sources or web-search context rather than relying on training weights.

Related on Tygart Media: how to use Claude · Anthropic API key.

Frequently Asked Questions

What is the knowledge cutoff for Claude Opus 4.7?

January 2026, per Anthropic's Help Center documentation verified May 15, 2026. Information about events, products, or developments after that date is not in the model's training data and must be provided directly.

What is the largest Claude context window currently available?

1 million tokens across all four current models (**Fable 5.2, Opus 5.5, Sonnet 5.5, and Haiku 5.5**), with 128K max output (legacy Haiku 4.5 is 200K context / 64K output).

Has Anthropic officially announced Claude 5?

Yes — starting in June 2026, Anthropic rolled out named 5.x families rather than one model named "Claude 5", advancing to **Fable 5.2, Opus 5.5, Sonnet 5.5, and Haiku 5.5** by November 2026.

Is Claude Sonnet 5 a real model I can use?

Yes — `claude-sonnet-5` shipped June 30, 2026 ($2/$10, 64K output) and is now superseded for new builds by **Claude Sonnet 5.5** (`claude-sonnet-5-5`, $2/$10, 1M context / 128K output).

Why does Opus 4.7 use more tokens than Opus 4.6 for the same text?

Opus 4.7 ships with a new tokenizer that contributes to its improved performance but uses approximately 1x to 1.35x as many tokens for the same input text compared to previous models. Anthropic recommends increasing `max_tokens` headroom and adjusting compaction triggers accordingly.

Are sampling parameters (temperature, top_p, top_k) still supported on Opus 4.7?

No. Setting `temperature`, `top_p`, or `top_k` to any non-default value on Opus 4.7 returns a 400 error. Migration guidance: omit these parameters and use prompting to guide the model's behavior.

Related Reading

- Claude Agent SDK Dual-Bucket Billing
- Claude MCP Token Cost Reality
- Amazon Prime Student + Claude Pro

How we sourced this

Sources reviewed May 15, 2026:

- Anthropic Pricing Documentation: *docs.claude.com/en/docs/about-claude/pricing* (primary source for model lineup, per-token rates, context window pricing)
- Anthropic Platform Documentation: *What's new in Claude Opus 4.7* (primary source for Opus 4.7 features, breaking changes, tokenizer, image support, task budgets)
- Anthropic Help Center: *How up-to-date is Claude's training data?* (primary source for knowledge cutoff dates)
- Anthropic news page (primary source check for Claude 5 announcement — none located as of May 15, 2026)
- Third-party reporting on Claude 5 / Sonnet 5 (TechCrunch interview reports, Claude5.com, Fello AI, WaveSpeed Blog) — cited as reported but not independently confirmed against primary sources

This article applies the verified vs. reported distinction throughout. If any of the unverified third-party claims are confirmed by Anthropic in the weeks after this article's date stamp, the relevant sections should be updated to reflect the new primary-source documentation.

Primary & Reference Links Cited on This Page

- Anthropic models overview — `https://platform.claude.com/docs/en/about-claude/models/overview`
- API pricing — `https://platform.claude.com/docs/en/about-claude/pricing`
- Anthropic news — `https://www.anthropic.com/news`
- Tygart Media: current model tracker
- Tygart Media: Claude Fable guide
- Tygart Media: how to use Claude
- Tygart Media: Anthropic API key
- Tygart Media: Claude Agent SDK Dual-Bucket Billing
- Tygart Media: Claude MCP Token Cost Reality
- Tygart Media: Amazon Prime Student + Claude Pro