TYGART MEDIA · CLAUDE & ANTHROPIC REFERENCE

What Changed in Claude This Month: November 2026

Edition of November 13, 2026 — facts verified November 13, 2026. Current edition always at https://tygartmedia.com/what-changed-in-claude-november-2026/.

Direct Answer (November 13, 2026): Anthropic's current November 2026 model lineup is led by **Claude Fable 5.2** (`claude-fable-5-2`, $8/$40, $0.20 cache read — down from Fable 5.1's $10/$50), alongside **Opus 5.5** ($4/$20), **Sonnet 5.5** ($2/$10), and **Haiku 5.5** ($0.10/$0.50 ≤100K). Plan naming finished settling on **Max ($100, 5× Pro + $100 API credits)** and **Ultra ($200, formerly Max 20x — 10× Pro weekly + no 5-hour cap for verified accounts + $200 API credits)**, Claude Code 2.1.317–2.1.319 hardened terminal permissions, and Anthropic's updated Usage Policy took effect November 12.

Last verified: November 13, 2026 · Primary sources: Anthropic models overview, API pricing & claude.com/pricing.

November started quietly on the model front before **Fable 5.2** took the top spot over Fable 5.1 ($8/$40 per MTok vs. $10/$50), while workflow changes hit every day: Claude got a read-only window into your bank accounts, Claude Code got quieter and safer in the terminal, the plan names finished shifting to Max and Ultra, and Anthropic's updated Usage Policy took effect November 12. The standing reference is our Claude model tracker, which stays current between editions.

Models

Fable 5.2 takes the top slot above Opus 5.5, Sonnet 5.5, and Haiku 5.5. An earlier November 8 draft of this digest noted a quiet week before the lineup updated — as of November 13, Anthropic's models overview and pricing schedule put **Fable 5.2** (`claude-fable-5-2`, $8/$40, $0.20 cache read) at the top, followed by **Opus 5.5** ($4/$20, $0.20 cache read), **Sonnet 5.5** ($2/$10, $0.10 cache read), and **Haiku 5.5** ($0.10/$0.50 ≤100K; $0.50/$2.50 >100K), all with a 1M-token context window and 128K max output. Fable 5.1 ($10/$50) moves to legacy. What we'd do: pin daily workloads to Opus 5.5 or Sonnet 5.5 and switch any Fable 5.1 callers to `claude-fable-5-2`.

One retirement is now close enough to act on: Sonnet 4.5, November 30. Anthropic deprecated Claude Sonnet 4.5 on September 30, and it retires November 30, 2026; the named replacement is Sonnet 5.5. If any integration, script, or saved workflow still calls `claude-sonnet-4-5-20250929`, this is the month it breaks. What we'd do: search your configs for that model ID this week, not on the 29th.

Plans & limits

The October 30 plan rename is the one still reverberating. What used to be Max 5x is now just Max ($100 a month, five times Pro's usage every five hours and every week), and Max 20x became Ultra ($200 a month, ten times Pro's weekly usage). Prices didn't change; the names did. The bigger deal inside the rename: verified Ultra accounts have **no five-hour session limit at all** — if your Settings > Usage screen shows no "Current session" bar on Ultra, that's not a bug, that's the tier working. October also started the monthly API credits that come with the paid tiers — $100 a month on Max, $200 on Ultra, up to $500/month pooled per Team org — claimable once you've been on the plan seven days by linking a Console organization in Settings > Billing. Plenty of people are paying for headroom they haven't claimed. The full map is our pricing, plans, and limits guide; the month this matters most is the month you hit the wall, and for that we wrote the playbook: You Hit Your Claude Limit. Now What?

November 5: Claude can now look at your money — read-only. Personal Finance in Claude launched as a beta for Pro, Max, and Ultra users in the US. You connect bank and card accounts through Plaid (you sign in at your bank; Claude never sees the password), and Claude can see balances, credit limits, and up to two years of transactions — enough to flag a duplicate charge, a fee you didn't notice, or a subscription that's been creeping up. It cannot move, manage, or invest money; it's a window, not a hand. Anthropic also says conversations in the finance hub aren't used for training and there are no ads or sponsored recommendations in it. What we'd do: if you connect it, connect one account first and see whether the answers earn the second one. Read-only is the right shape for a first version of this.

November 12: the updated Usage Policy takes effect. Anthropic published the revision October 8; it becomes the rules this month. The plain-language version: there's a new consolidated section against deceptive campaigns and artificial

activity — fake accounts, fabricated news sites, hiding who's behind a message. The elections section was rewritten and its blanket ban on personalized campaign targeting was *removed* (deception-based targeting is still banned). The surveillance rules got sharper: tracking people without consent is out, and Claude can't be used to decide whom to investigate, arrest, or charge. There's also a new line against sustained, needless cruelty toward the models themselves — Anthropic is explicit that ordinary frustration, pushback, and testing don't count. What we'd do: if your shop publishes or campaigns with Claude in the loop, read the deceptive-campaigns section once, in full. It's the section most likely to describe something a marketing workflow does by accident.

Claude Code

Three releases landed in the first week of November (2.1.317 through 2.1.319), and the theme is the terminal getting more careful about what leaves your machine.

November 3 (2.1.317): Claude Code now asks before a web or shell tool call sends your account email address out to the network. Small change, correct instinct — your email is exactly the kind of thing that leaks by accident inside an agent run. The `[1m]` suffix also disappeared from model names and IDs, so old screenshots and scripts that reference it are describing a naming scheme that's gone.

November 4 (2.1.318): Your sign-in now tells you when it expires. `claude auth status` shows a `loginExpiresAt` date, scripted runs surface `credential_expires_at`, and the terminal starts warning you ten days before expiry instead of letting a session die mid-job. Cloud sessions can also start with no repository attached. Anyone who has had an overnight run fail on a dead login will read this item twice.

November 6 (2.1.319): The busy one. MCP tasks in the terminal now move to the background and report back when they're done instead of holding your session hostage. There's a new `permissions.denyInsteadOfPrompting` setting — deny outright instead of asking — and auto mode no longer falls back to a built-in classifier when its server check can't run; it asks or denies (a temporary environment variable restores the old behavior; Bedrock, Vertex, and Foundry are unchanged). Ninety bug fixes ride along, including a worktree sandbox-escape fix and a Bash permission-check bypass fix — the kind of fixes you want promptly and read about calmly.

Also November 6: cloud sessions learned to set themselves up. On Pro, Max, and Ultra, Claude Code cloud sessions now work out a repository's dependencies, network access, and API keys on their own and save the environment for next time. The detail that matters: keys are stored as "network secrets" and injected by Anthropic's agent proxy from outside the session machine, so the model never sees them. Setup was the tax on every cloud session; a chunk of that tax just went away. And one leftover from October 31 with a November deadline: fast mode for Opus 4.8 was deprecated October 30 and is **removed November 30** — after that, fast-mode requests on Opus 4.8 error instead of falling back to standard speed. If you want the speedup, the migration target is Opus 5.5.

API & Console

November's API news is mostly the absence of drama — plus two October changes whose consequences land now. The monthly API credits described above (started October 7) are live money for Max, Ultra, and Team accounts, but they come with a fence readers keep tripping on: they cover the Claude API, Managed Agents, the Agent SDK, and scripted `claude -p` runs with an API key — they do **not** cover interactive Claude Code or extra usage inside the apps. Two different buckets, two different rules; our Console guide covers the account side. And for enterprise shops: an October-announced security package — inference hooks that route prompts and tool calls through your own security stack before anything runs, a Security Center, and endpoint detection inside Claude's unified desktop/agent environment (Anthropic folded standalone Cowork into core Claude — see claude.com/product/cowork) — is rolling out through the fall, Enterprise tier only. Most small shops can file that under "good to know exists" and move on.

The short version

November updated both the flagship tier and the daily ergonomics: Fable 5.2 ($8/$40) at the top, one retirement to beat (Sonnet 4.5, November 30), plan names and perks from late October still settling, Claude Code getting more careful and more self-sufficient, and a policy update that takes effect November 12. The living reference for all of it is the model tracker — this digest runs monthly, and December's edition lands when there's a month of changes worth your coffee.

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`
- claude.com/pricing — `https://claude.com/pricing`
- claude.com/product/cowork — `https://claude.com/product/cowork`
- Tygart Media: Claude model tracker
- Tygart Media: pricing, plans, and limits guide
- Tygart Media: You Hit Your Claude Limit. Now What?