Claude Code Pricing 2026: Pro, Max, and API Limits

Edition of November 13, 2026 — facts verified November 13, 2026. Current edition always at https://tygartmedia.com/claude-code-pricing-2026/.

Direct Answer (November 13, 2026): Claude Code runs on **Pro ($20/month or $17/month billed annually)**, **Max ($100/month, 5× Pro usage + $100/mo API credits)**, **Ultra ($200/month, formerly Max 20× — 10× Pro weekly usage, no 5-hour limit for verified accounts + $200/mo API credits)**, **Team ($20–$25/seat Standard or $100–$125/seat Premium with 5× Standard usage, 2–150 seats)**, **Enterprise ($20/seat/mo annual + API usage, 20-seat min)**, or pay-as-you-go API tokens (**Fable 5.2 $8/$40**, **Opus 5.5 $4/$20**, **Sonnet 5.5 $2/$10**, **Haiku 5.5 $0.10/$0.50** per MTok ≤100K).

Updated October 28, 2026: Made infographic diagrams clickable to open full-size images.

Last refreshed: November 13, 2026 (Pacific)

Last verified: November 13, 2026 against claude.com/pricing, Claude API pricing, and Anthropic's Ultra plan Help Center docs.

Claude Code pricing has stopped being a clean sticker number and started being a question of which ceiling you hit first. There is a $20 plan, a $100 plan, and a $200 plan — and underneath all three sits a 5-hour rolling window, a weekly active-hours cap added in August 2025, and a per-model multiplier that quietly makes Opus 5.5 the most expensive thing you can do inside the terminal. If you came looking for the right plan, the honest answer is: it depends on whether you are mostly a Sonnet operator or you live in Opus.

For the full picture across all plans, see the Claude pricing overview.

The three subscription tiers, stripped down

The three subscription tiers, stripped down.

Pro — $20/month ($17/month billed annually at $200 upfront). Access to Claude Code in the terminal, web, and desktop, with Sonnet 5.5 and Opus 5.5 available ($0 API credits included; Fable models bill via usage credits). Pro provides at least 5× Free usage per 5-hour session plus a weekly limit. This is the plan for occasional coding sessions, code review, and moderate refactoring where you can tolerate hitting a reset clock mid-afternoon.

Max — $100/month (monthly only). Published as **5× Pro usage** every 5 hours and weekly, plus included Fable access up to 50% of the weekly limit and **$100/month in Claude Platform API credits**. If you spend most of your working week inside Claude Code, Max is where the 5-hour ceiling stops interrupting normal daily sprints.

Ultra — $200/month (monthly only; formerly Max 20×). Anthropic replaced Max 20× with Ultra at the same $200/month price: published as **10× Pro weekly usage**, **no 5-hour session limit for verified accounts** (temporary 5-hour limits apply if unverified), Fable access up to 50% of weekly limits, and **$200/month in Claude Platform API credits**. Heavy multi-hour autonomous runs can still hit the weekly ceiling, at which point pay-as-you-go API usage or usage credits cover overflow.

The API, as a sanity check

The API, as a sanity check.

If you want a sanity check on whether the subscription math works, price the same workload against the API:

- **Claude Haiku 5.5** (`claude-haiku-5-5`): $0.10 input / $0.50 output per MTok (prompts ≤100K tokens; $0.50/$2.50 above 100K; legacy Haiku 4.5 is $1.00/$5.00 with 200K context)
- **Claude Sonnet 5.5** (`claude-sonnet-5-5`): $2.00 input / $10.00 output per MTok ($0.10 cache read)
- **Claude Opus 5.5** (`claude-opus-5-5`): $4.00 input / $20.00 output per MTok ($0.20 cache read)
- **Claude Fable 5.2** (`claude-fable-5-2`): $8.00 input / $40.00 output per MTok ($0.20 cache read; current flagship, superseding Fable 5.1 at $10/$50)

Prompt caching is the lever almost nobody uses correctly. Cache writes cost 1.25x input price for the 5-minute TTL or 2.0x for the 1-hour TTL, but cache reads cost 0.05x — a 95% discount on every subsequent request that hits the same context. If your `.clauderules` file, project map, and the file you are editing are all stable for an hour, the bill on a long pairing session can drop by an order of magnitude. The Batch API knocks another 50% off both directions for asynchronous workloads, which is worth knowing if you are running large refactor sweeps.

One trap on Opus 4.7 specifically: the model uses a new tokenizer that inflates token counts by up to 35% on identical text compared to Opus 4.6. The headline price did not change, but your effective spend per request did — sometimes by nothing, sometimes by a third, depending on the content. If you migrated from Opus 4.6 and your bill went up without your prompt patterns changing, that is the reason.

How to actually choose

How to actually choose.

The cleanest way to pick a plan is to first decide your model mix, then your weekly hours.

If you are mostly a **Sonnet operator** — long agentic runs, multi-file edits, codebase Q&A, with Opus only reached for on the architectural questions — Pro at $20 is plausible for focused daily work, Max ($100/mo, 5× Pro) covers most full-time individual developers, and Ultra ($200/mo, 10× Pro weekly + no 5-hour limit for verified accounts) covers heavy parallel sessions.

If you live in **Opus** — long-horizon agentic work, hard refactors across many files, anything where you would rather have one good attempt than three Sonnet retries — Pro will frustrate you quickly, Max ($100/mo) is the realistic floor, and Ultra ($200/mo) removes the 5-hour session cap for verified accounts so you can run multi-hour tasks without bouncing to API billing.

And if you are running Claude Code **across multiple repos all week**, leaving agents to grind on tasks while you do other things, Ultra ($200/mo) is the highest individual subscription tier — and even then, the weekly cap is real. Use the API for the spillover and you will still come out cheaper than trying to brute-force a smaller plan.

The number that matters

One developer's public report this year: roughly 10 billion tokens consumed across Claude Code over eight months. API metered cost would have exceeded $15,000. The same workload on Max at $100/month for the same window came in around $800 — about 93% cheaper. That is the gap that makes the subscription model worth taking seriously, even when the rate limits feel arbitrary. The $200 tier is not a vanity number; it is the price Anthropic charges to stop being a meaningful constraint on your workflow.

The right way to read Claude Code pricing in November 2026 is not to ask which plan is cheapest. It is to ask which plan is the cheapest one that disappears — the one that stops appearing in your day. For most full-time developers reaching for Opus or Fable regularly, that plan is Max ($100/mo) or Ultra ($200/mo). For everyone else, start at Pro ($20/mo or $17/mo annual) and watch Settings → Usage before upgrading.

■ Recommended Reading in Claude Code Insider

- ■ Pillar Guide: Claude Code + GitHub in 2026: What Rakuten, TELUS, and a 100K-Star Config File Actually Reveal
- ■ Next Topic: Claude Code Case Studies: What the Numbers Actually Say in 2026

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

Two meters

- **Seat** — claude.ai / Claude Code usage cap on Pro, Max, Team, Enterprise.
- **API** — prepaid credits on platform.claude.com. Same models, different bill.

Official seats: claude.com/pricing and the Team Help Center article. Official tokens: API pricing.

Token rates used inside Code (API path)

Model	In / out per MTok	When
Haiku 5.5	$0.10 / $0.50 (≤100K) · $0.50 / $2.50 (>100K)	Cheap lookups
Sonnet 5.5	$2 / $10	Current Sonnet
Opus 5.5	$4 / $20	Current Opus. Docs say start here
Opus 5	$5 / $25	Legacy. Still listed
Fable 5.2	$8 / $40	Only if the job is worth it

Fable 5.1 ($10 / $50), Haiku 4.5 ($1 / $5, 200K context), Sonnet 4.6 ($3 / $15), Opus 4.7 ($5 / $25), and Opus 4.8 ($5 / $25) remain on the legacy list — start new work on Fable 5.2, Opus 5.5, Sonnet 5.5, or Haiku 5.5. Fast mode runs at 2× standard pricing on Fable 5.2, Opus 5.5, Opus 5, and Opus 4.8.

Which seat

- Light Code: Pro.
- Several hours a day: Max ($100/mo, 5× Pro).
- All-day / long agents: Ultra ($200/mo, formerly Max 20× — 10× Pro weekly + no 5-hour limit for verified accounts).
- 2–150 people, one bill: Team. Premium if those seats burn the weekly cap.
- SSO / SCIM / audit and usage that should scale: Enterprise.
- Unattended pipelines: platform.claude.com API key with a spend cap.

Related: current models · console setup.

Primary & Reference Links Cited on This Page

- claude.com/pricing — `https://claude.com/pricing`
- Claude API pricing — `https://platform.claude.com/docs/en/about-claude/pricing`
- Ultra plan Help Center docs — `https://support.claude.com/en/articles/17728952-what-is-the-ultra-plan`
- Tygart Media: Claude pricing overview
- Tygart Media: https://tygartmedia.com/wp-content/uploads/2026/08/claude-tiers-at-a-glance-board-1.jpg
- Tygart Media: https://tygartmedia.com/wp-content/uploads/2026/08/claude-api-token-meter-visual.jpg
- Tygart Media: https://tygartmedia.com/wp-content/uploads/2026/08/claude-which-plan-decision-map.jpg
- Tygart Media: Claude Code + GitHub in 2026: What Rakuten, TELUS, and a 100K-Star Config File Actually Reveal
- Tygart Media: Claude Code Case Studies: What the Numbers Actually Say in 2026
- Tygart Media: how to use Claude
- Tygart Media: Anthropic API key
- Tygart Media: current models
- Tygart Media: console setup