Grok vs Claude: Seats, API Rates, and the Real Cost Gap

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

Direct Answer (November 13, 2026): At the API entry tier, **Claude Haiku 5.5 ($0.10/$0.50 per MTok ≤100K, 1M context)** now undercuts Grok Build 0.1 ($1/$2) and Grok 4.3 ($1.25/$2.50), while at the flagship tier Grok 4.7 ($2/$6) remains cheaper on list tokens than Claude **Opus 5.5 ($4/$20)** and **Fable 5.2 ($8/$40, down from Fable 5.1's $10/$50)**. On consumer plans, Claude Pro ($20/mo) and Ultra ($200/mo, 10× Pro weekly, no 5h cap for verified accounts) sit below SuperGrok (~$30/mo) and Heavy (~$300/mo), with Fable included up to 50% of weekly usage on Max ($100), Ultra ($200), and Team Premium / Enterprise seats.

Last verified: November 13, 2026 (Pacific) against claude.com/pricing, Anthropic API pricing, Anthropic's Ultra guide, and published xAI developer pricing. Consumer seat prices vary by store and region — confirm before you pay.

Consumer seats

	Grok (xAI)	Claude (Anthropic)
Free	Metered on grok.com and X	Metered, Sonnet
Everyday paid	SuperGrok ~$30/mo (X Premium+ bundle ~$40)	Pro $20/mo
Heavy individual	Plus ~$100 or Heavy ~$300	Max $100/mo (5x Pro usage); Ultra $200/mo (10x Pro weekly)
Team	Grok Business ~$30/seat	Team Standard $20/seat annual ($25 monthly); Premium $100 annual / $125 monthly

Claude Pro ($20/mo; $17 annual) wins the entry consumer sticker, and Ultra ($200/mo, formerly Max 20x — 10× Pro weekly with no 5-hour session limit for verified accounts + $200/mo API credits) undercuts SuperGrok Heavy ($300). Claude gates its top tier by plan: **Fable 5.2** is included on Max ($100/mo, $100 API credits), Ultra ($200/mo), and Team Premium / Enterprise seats (up to 50% of the weekly limit per Anthropic's seat rules). Pro ($0 API credits) and Team Standard ($20/$25/seat, includes Claude Code + up to $500/mo pooled org API credits) draw usage credits when calling Fable. Details: Fable pricing and plan access and Claude Code limits.

API rates per million tokens

Grok. Official xAI card. Prompts that reach 200k tokens are billed at 2x for the whole request.

Model	In / out	Context	Cached input
Grok Build 0.1	$1 / $2	256k	$0.20
Grok 4.3 / 4.20	$1.25 / $2.50	1M	$0.20
Grok 4.5 / 4.6 / 4.7	$2 / $6	500k	$0.30-$0.50

Claude.

Model	In / out	Context	Cache read
Haiku 5.5 (current)	$0.10 / $0.50 (≤100K); $0.50 / $2.50 (>100K)	1M	$0.01 / $0.05
Haiku 4.5 (legacy)	$1 / $5	200k	$0.10
Sonnet 5.5 (current)	$2 / $10	1M	$0.10
Opus 5.5 (current)	$4 / $20	1M	$0.20
Fable 5.2 (current flagship)	$8 / $40	1M	$0.20

Model	In / out	Context	Cache read
Fable 5.1 (legacy)	$10 / $50	1M	$0.25

Fable 5.2 reduced Anthropic's flagship list price from $10/$50 (Fable 5.1) to **$8/$40 per MTok**, with cache reads at **$0.20** (down from $0.25), 5m/1h cache writes at $10/$16, and Batch at $4/$20. Every current 5.x Claude model (including Haiku 5.5) now carries a 1M context window and 128K max output.

Same-class pairing

- Budget / high-volume: **Haiku 5.5 ($0.10/$0.50 ≤100K)** vs Grok Build 0.1 ($1/$2). Claude wins decisively on short prompts under 100K tokens.
- Everyday production: Grok 4.3 ($1.25/$2.50) vs Sonnet 5.5 ($2/$10). Grok is cheaper, especially on output.
- Flagship work: Grok 4.7 ($2/$6) vs Opus 5.5 ($4/$20). Grok remains cheaper on list tokens. (Legacy Opus 5 remains $5/$25.)
- Top shelf: Grok has no $8/$40 frontier SKU. Fable 5.2 is 4× Grok 4.7 input and ~6.7× output.

Worked example

10M input + 2M output, no cache, short prompts (≤100K):

- Haiku 5.5 (≤100K): $1 + $1 = **$2**
- Grok 4.6 / 4.7: $20 + $12 = **$32**
- Sonnet 5.5: $20 + $20 = **$40**
- Opus 5.5: $40 + $40 = **$80**
- Fable 5.2 (current): $80 + $80 = **$160** (down from $200 on legacy Fable 5.1)

Batch: Claude 50% off on supported models (Fable 5.2 $4/$20, Opus 5.5 $2/$10, Sonnet 5.5 $1/$5, Haiku 5.5 $0.05/$0.25 ≤100K). Grok 20% off on 4.3 / 4.20 only – not on 4.5 / 4.6.

Rules that are not the rate card

- Claude individual plans use a rolling 5-hour session cap and a weekly bucket — except **Ultra ($200/mo, 10× Pro weekly)**, which removes the 5-hour limit for verified accounts; **Max ($100/mo)** sits at 5× Pro, and **Team Premium ($100/$125)** provides 5× Team Standard usage.
- Grok API doubles the request once the prompt hits 200k tokens. Current Claude 5.x cards include the full 1M window at standard per-token rates (Haiku 5.5 uses a $0.50/$2.50 tier above 100K prompts).
- Prompt caching keeps top-tier runs tractable: Fable 5.2 cache reads cost **$0.20/MTok** (vs. $8/$40 fresh prompts); Opus 5.5 cache reads are $0.20; Sonnet 5.5 cache reads are $0.10.

When to buy which

Buy Grok for volume, agents, or coding at $2/$6 where Grok 4.6 is enough.

Buy Claude when the job needs Fable-class long horizon and you will pay for it – or when Sonnet 5.5 at $2/$10 is enough and the team already lives in Claude Code.

Do not pick from the consumer sticker alone. A $20 Claude Pro seat that immediately burns Fable credits can cost more than a $30 SuperGrok seat that never leaves Grok 4.6.

Claude vs Grok: which is better for real-time information?

When evaluating real-time information retrieval, Grok and Claude operate on fundamentally different ingestion models as officially documented by xAI and Anthropic:

- **Grok's real-time retrieval model:** Built by xAI, Grok is natively embedded within the X (formerly Twitter) platform. It has continuous, low-latency access to real-time social posts, breaking global news, emerging trend conversations, and live

social sentiment. For events unfolding minute-by-minute—such as live breaking news, market-moving announcements, tech releases, or public sentiment monitoring—Grok provides immediate situational awareness that traditional search indexing cannot match.

- **Claude's web search retrieval model:** Anthropic provides real-time information through integrated web search capabilities across claude.ai and the Claude Platform API (priced at $10 per 1,000 web searches, separate from standard token rates). Rather than indexing social conversational firehoses, Claude's search tool queries authoritative web indices, retrieves live documentation, extracts factual sources, and filters results through Claude's rigorous safety and synthesis reasoning layers.
- **The practical split:** Choose **Grok** when your query demands up-to-the-minute freshness, social sentiment analysis, breaking news commentary, or monitoring live community reactions. Choose **Claude** when your real-time need involves factual verification, technical documentation research, synthesizing regulatory or corporate updates from published web sources, and drafting balanced analytical reports where depth and source grounding matter more than social media immediacy.

FAQ

Is SuperGrok cheaper than Claude Pro? No on the monthly line: Pro is $20, SuperGrok is about $30. Yes on many flagship API workloads, where Grok 4.7 ($2/$6) undercuts Claude Opus 5.5 ($4/$20) and Fable 5.2 ($8/$40).

Is Claude always more expensive on the API? No. Claude Haiku 5.5 ($0.10/$0.50 ≤100K) undercuts Grok Build 0.1 ($1/$2) and Grok 4.3 ($1.25/$2.50). The Claude premium starts at Opus 5.5 ($4/$20) and Fable 5.2 ($8/$40).

Does Claude Pro include Fable 5.2? No. Pro pays pay-as-you-go usage credits for Fable. Included Fable is on Max ($100), Ultra ($200), and Team Premium / Enterprise seats (50% of weekly limits on included plans). See Fable plan access.

Related: Claude plan pricing · Grok vs Claude (capability comparison – older lineup) · Claude Code limits.

If you run a restoration or multi-site operation and want the same kind of defensible, versioned standard for your Scope 3 emissions data, see the Restoration Carbon Protocol – the open framework that maps contractor emissions onto the GHG Protocol so commercial clients can actually verify them.

Primary & Reference Links Cited on This Page

- claude.com/pricing — `https://claude.com/pricing`
- Anthropic API pricing — `https://platform.claude.com/docs/en/about-claude/pricing`
- Anthropic's Ultra guide — `https://support.claude.com/en/articles/17728952-what-is-the-ultra-plan`
- Tygart Media: Fable pricing and plan access
- Tygart Media: Claude Code limits
- Tygart Media: Claude plan pricing
- Tygart Media: Grok vs Claude (capability comparison – older lineup)
- Tygart Media: Restoration Carbon Protocol