TYGART MEDIA · CLAUDE & ANTHROPIC REFERENCE

How Many Words Is a Million Claude Tokens? (2026)

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

Direct Answer (November 13, 2026): All four current Claude models — **Fable 5.2 ($8/$40), Opus 5.5 ($4/$20), Sonnet 5.5 ($2/$10), and Haiku 5.5 ($0.10/$0.50 ≤100K)** — have a **1M-token context window** and 128K max output (models overview). Using Anthropic's newer Opus 4.7+ tokenizer (~1.30–1.35x tokens vs. older Claude 4.6 models), 1M tokens holds ~555,000 English words (~2,200 pages), compared with ~750,000 words (~3,000 pages) under the legacy 4.6 tokenizer; legacy Haiku 4.5 is capped at 200K tokens.

Last verified: November 13, 2026 against Anthropic models overview and Token counting docs.

A million Claude tokens equals roughly **750,000 words** on Claude Sonnet 4.6 â€" but only about **555,000 words** on Claude Opus 4.7, Claude Opus 4.8, and Claude Fable 5. The gap comes from a new tokenizer that Anthropic introduced with Opus 4.7: it emits up to 35% more tokens from the same text. The only reliable way to measure your actual token count is the `/v1/messages/count_tokens` endpoint.

Token-to-word conversion by model (1 million tokens)

Token-to-word conversion framing.

Anthropic publishes word equivalents directly in the context-window tooltips on the official models overview page. The figures below come from those tooltips.

Lineup currency (November 13, 2026): Current API lineup (all 1M context / 128K output): Fable 5.2 $8/$40, Opus 5.5 $4/$20, Sonnet 5.5 $2/$10, Haiku 5.5 $0.10/$0.50 (≤100K; $0.50/$2.50 >100K). Legacy models in the tokenizer comparison tables below (Fable 5, Opus 4.8, Opus 4.7, Sonnet 4.6, Haiku 4.5 at 200K context) are preserved specifically to show the pre- vs. post-Opus 4.7 tokenizer split.

Model	Tokenizer	Context window	~Words per 1M tokens	~Pages per 1M tokens*
Claude Fable 5 (`claude-fable-5`)	New (Opus 4.7)	1M tokens	~555,000	~2,200
Claude Opus 4.8 (`claude-opus-4-8`)	New (Opus 4.7)	1M tokens	~555,000	~2,200
Claude Opus 4.7 (`claude-opus-4-7`)	New (Opus 4.7)	1M tokens	~555,000	~2,200
Claude Sonnet 4.6 (`claude-sonnet-4-6`)	Older	1M tokens	~750,000	~3,000
Claude Haiku 4.5 (`claude-haiku-4-5`)	Older	200k tokens	~150,000 (200K context)	~600 (200K context)
Claude Opus 4.6 (`claude-opus-4-6`)	Older	1M tokens	~750,000	~3,000

* Pages estimated at ~250 words per double-spaced page. These are approximations for typical English prose; actual counts vary by content type.

What the new tokenizer changed â€" and why it matters

What the new tokenizer changed — and why it matters.

Anthropic introduced a new tokenizer with Claude Opus 4.7. The official migration guide states that the new tokenizer "may use roughly 1x to 1.35x as many tokens when processing text compared to previous models (up to ~35% more, varying by content)." The most commonly cited figure across Anthropic's documentation is roughly 30% more tokens for the same text.

The practical effect: a document that costs 1,000,000 tokens on Opus 4.6 or Sonnet 4.6 costs approximately 1,300,000 tokens on Opus 4.7, Opus 4.8, or Fable 5. Budgets built for the old tokenizer need to be re-baselined against the new one.

Lineup currency (November 13, 2026): Current API lineup (all 1M context / 128K output): Fable 5.2 $8/$40, Opus 5.5 $4/$20, Sonnet 5.5 $2/$10, Haiku 5.5 $0.10/$0.50 (≤100K; $0.50/$2.50 >100K). Legacy models in the tokenizer comparison tables below (Fable 5, Opus 4.8, Opus 4.7, Sonnet 4.6, Haiku 4.5 at 200K context) are preserved specifically to show the pre- vs. post-Opus 4.7 tokenizer split.

Tokenizer	Models	Approximate token increase vs. older tokenizer
New (introduced Opus 4.7)	Opus 4.7, Opus 4.8, Fable 5, Mythos 5	~30% typical; up to ~35% depending on content
Older	Opus 4.6, Sonnet 4.6, Haiku 4.5, Opus 4.5, Sonnet 4.5	Baseline

The token counting page also notes the comparison directly: "Claude Fable 5 and Claude Mythos 5 use the tokenizer introduced with Claude Opus 4.7, which produces roughly 30% more tokens than models before Claude Opus 4.7 for the same text."

Use count_tokens â€” not tiktoken or ratio math

Use count_tokens — not tiktoken or ratio math.

Anthropic's migration guide explicitly flags the risk: "Any code path that estimates tokens client-side or assumes a fixed token-to-character ratio should be re-tested against Claude Opus 4.7." OpenAI's tiktoken library is trained on a different vocabulary and produces different counts. It will not give accurate results for any Claude model.

The correct approach is the `/v1/messages/count_tokens` endpoint, passing the specific model you intend to use:

```
curl https://api.anthropic.com/v1/messages/count_tokens \
  --header "x-api-key: $ANTHROPIC_API_KEY" \
  --header "content-type: application/json" \
  --header "anthropic-version: 2023-06-01" \
  --data '{
    "model": "claude-opus-5-5",
    "messages": [{"role": "user", "content": "Your text here"}]
  }'
```

The endpoint returns a model-specific count. If you are migrating a workload from a legacy model like Sonnet 4.6 to a current model like Opus 5.5 or Fable 5.2, count the same prompt with both model IDs and compare the two `input_tokens` values. The token counting endpoint is free to use (rate limits apply by usage tier). Anthropic notes that the returned count is an estimate; the actual count at inference time may differ by a small amount.

Quick reference: common document sizes

Lineup currency (November 13, 2026): Current API lineup (all 1M context / 128K output): Fable 5.2 $8/$40, Opus 5.5 $4/$20, Sonnet 5.5 $2/$10, Haiku 5.5 $0.10/$0.50 (≤100K; $0.50/$2.50 >100K). Legacy models in the tokenizer comparison tables below (Fable 5, Opus 4.8, Opus 4.7, Sonnet 4.6, Haiku 4.5 at 200K context) are preserved specifically to show the pre- vs. post-Opus 4.7 tokenizer split.

Document type	Approx. words	Tokens (older tokenizer)	Tokens (new tokenizer)
Novel (~400 pages)	~100,000	~133,000	~173,000
Long research paper	~20,000	~27,000	~35,000
Full context, Sonnet 4.6 (1M tokens)	~750,000	1,000,000	N/A (different model)
Full context, Opus 4.8 (1M tokens)	~555,000	N/A (different model)	1,000,000

These word estimates assume typical English prose. Code, structured data, and non-Latin scripts tokenize differently from natural language prose. Highly repetitive text and dense symbol-heavy content (like JSON or code) can fall well outside the ~0.75 words-per-token ratio.

Does the new tokenizer change what fits in the context window?

Yes, in one direction. The context window is still 1M tokens, but that window holds fewer words on the new tokenizer (~555k words) than on the old one (~750k words). A document that previously fit comfortably may now require trimming or chunking when moving to Opus 4.7, Opus 4.8, or Fable 5.

Does Sonnet 4.6 use the new tokenizer?

No. Claude Sonnet 4.6 uses the older tokenizer. Anthropic's model overview page lists Sonnet 4.6's 1M-token context window as equivalent to ~750k words, the same ratio as Opus 4.6 â€” confirming it has not adopted the Opus 4.7 tokenizer. Only Opus 4.7, Opus 4.8, Fable 5, and Mythos 5 use the new tokenizer.

Can I use tiktoken or another open-source tokenizer for Claude?

No. tiktoken is built for OpenAI models and uses a different vocabulary. It will not produce accurate token counts for any Claude model, and its error will be larger on the new Opus 4.7 tokenizer than on older Claude models. Use `/v1/messages/count_tokens` with the specific Claude model ID you plan to deploy.

Does the new tokenizer affect pricing?

Yes. Billing reflects token counts under the model's tokenizer. If you migrate a workload from Opus 4.6 to Opus 4.8 and the new tokenizer produces 30% more tokens, your input token costs increase by roughly 30% before accounting for any per-token price difference between the models. Re-baseline cost estimates using the count_tokens endpoint rather than scaling from old measurements.

How many pages is the full 1M-token context window?

On models with the older tokenizer (Sonnet 4.6, Opus 4.6), 1 million tokens is approximately 3,000 double-spaced pages of typical English prose. On models with the new tokenizer (Opus 4.8, Fable 5), the same 1 million tokens holds approximately 2,200 pages. These are prose estimates â€” a 1M-token window filled with source code or dense structured data will span a very different page count.

Related on Tygart Media: Message Batches API · Claude pricing · Anthropic API quickstart.

Primary & Reference Links Cited on This Page

- Anthropic models overview — `https://platform.claude.com/docs/en/about-claude/models/overview`
- Token counting docs — `https://platform.claude.com/docs/en/build-with-claude/token-counting`
- API lineup — `https://platform.claude.com/docs/en/about-claude/pricing`
- Tygart Media: Message Batches API
- Tygart Media: Claude pricing
- Tygart Media: Anthropic API quickstart