



— AGENTBURN · INDUSTRY REPORT · Q4 2026

The Agent Token Economy

What 30 production AI agents actually burn: tokens per task, cost per solved task, and framework overhead — measured against public benchmarks and priced from official model rates.

Dataset snapshot: 2026-10-21 · 30 projects · 22 models · 3 benchmarks
Prepared by the AgentBurn editorial desk · getagentburn.com
This edition is free to cite with attribution.



getagentburn.com/report · The Agent Token Economy · Q4 2026

What is inside

- 01 Executive summary
- 02 Methodology and provenance
- 03 The fleet at a glance
- 04 Where the tokens go
- 05 Cost per solved task
- 06 Framework overhead
- 07 Model pricing
- 08 Public benchmarks
- 09 Retry overhead and context
- 10 What to watch next
- A Project profiles (30)
- B Cost vs success map
- C Build vs buy ROI matrix
- D Appendix: fields, tiers and citation

How to read this report. Every figure is traceable to the dataset snapshot. Provenance badges tell you whether a number is Verified against a public source, Community-reported, or Estimated from adjacent public data. Nothing here is audited billing data.

The bill is real, and it is concentrated

1.98T

tokens per day across the tracked fleet (modelled index)

89.9%

of tracked burn comes from coding agents

\$0.01 → \$35.71

cost per solved task, 3,143× spread

13 / 30

projects with a public source (Verified tier)

Agent spend is usually discussed per million tokens, which hides the only number a buyer can compare: the cost of a *solved* task. Once failures and retries are priced in, the fleet spans three orders of magnitude — from about a cent for a research task to more than \$35 for a single paper.

#	AGENT	DOMAIN	DAILY TOKENS	COST / SOLVED TASK
1	Claude Code	coding	831.00B	\$19.71
2	GitHub Copilot coding agent	coding	265.20B	\$1.15
3	Cursor Agent	coding	175.00B	\$0.15
4	OpenHands	coding	156.00B	\$1.88
5	Codex CLI	coding	144.00B	\$1.42

Top five by modelled daily token burn. Cost per solved task = cost per task ÷ success rate, with retries already inside the bill.

Takeaway. Cheap tokens are not cheap outcomes. The cheapest solved task in the dataset costs about 0.3% of the most expensive one, and the gap is driven as much by retry behaviour and context strategy as by list price.