

AI Currency Card — Aug 9, 2026

Podium reference. If a number isn't on this card, don't say it on stage. *Verified 2026-07-21 against the live-search-checked digest (07-15) + daily scans through 07-21. Re-verify at T-7 (Aug 2) and T-2 (Aug 7).*

The lineup (safe to name on stage)

Provider	Say this	Price (per M tokens in/out)	One-liner
Anthropic	Claude Fable 5 (frontier) · Opus 4.8 · Sonnet 5 · Haiku 4.5	\$10/\$50 · \$5/\$25 · \$2/\$10 intro → \$3/\$15 after Aug 31 · \$1/\$5	Sonnet 5 is the default on Free + Pro — the audience follows along at no cost
OpenAI	GPT-5.6 — Sol / Terra / Luna (Jul 9)	\$5/\$30 · \$2.50/\$15 · \$1/\$6	Consumer ChatGPT default is still GPT-5.5 Instant; 5.6 is the flagship/agentic tier — both true at once
xAI	Grok 4.5 (Jul 8)	\$2/\$6	"Opus-class," 500K context, aggressively cheap
Google	Gemini 3.5 Flash	Flash-tier	3.5 Pro is delayed (Bloomberg Jul 16) — Google trails the frontier right now

NEVER say as current: GPT-5.5 · Grok 4.20 · Sonnet 4.6 · Opus 4.5. **Don't pitch:** Claude Mythos 5 (invitation-only; not an audience option).

The July story (the narrative upgrade)

- | "In July, AI stopped being a chatbot and became a coworker."
- **Jul 9 — OpenAI wave:** ChatGPT **Work** (hand it a job, get finished docs/sheets/decks back) + unified **Desktop app** (free, Mac+Win) + **Scheduled Tasks**.
 - **Jul 8 — GPT-Live-1** real-time voice (interruptible, natural turn-taking). Voice-AI receptionists stopped being sci-fi.

- **Jul 16 — Grok Automations:** scheduled + email-triggered jobs, no code.
- **Claude:** scheduled routines + Skills + MCP connectors (what Session A teaches).
- **Implication for Deck B Part 2:** scheduled agentic jobs are now **no-code on all three platforms**. “The Mac-and-cron version I run is the power-user path — you can start with a toggle.”

Price story

Frontier-adjacent AI now runs \$1–6 per million tokens (Luna \$1/\$6, Grok \$2/\$6, Sonnet 5 \$2/\$10). **”Cost is no longer the barrier. Knowing what to automate is.”** Flag: Sonnet 5 intro pricing ends Aug 31 (3 weeks after today).

Compliance delta (feed to Zach, one-line callback in A3)

- **Anthropic self-serve HIPAA (Jul 14):** enterprise admin can review the BAA + enable HIPAA features without a sales call (Enterprise, ~50-seat min). New market baseline — “the vendors are meeting you halfway.”
- Texas SB 1188 (AI-diagnostic disclosure + US data localization) — in force, unchanged. Feb 16 2026 HIPAA update — in force. Federal Security Rule final — still pending.

Dental-AI proof points (with confidence labels)

-  **Pearl Oral Health Index** — 26M x-rays / 15M patients: cavity burden **4.5x worse than CDC estimates**. Strongest stat of the year. (Multiple scans, consistent.)
-  **Overjet** — 10th FDA clearance (Caries Assist, Jul 7); major DSO rollouts.
-  “1 in 3 US practices use AI” — watch-report-sourced; say “roughly a third by industry estimates.”
-  Funding/valuation figures (Pearl \$58M Series B, Overjet totals, Videa) — reports disagree; avoid precise dollars on stage.
-  **PatientDesk (YC W26)** — front-desk AI going \$17K→\$50K MRR in 8 weeks: “the ops layer is where dental AI is moving — exactly what you’re learning today.”

Scan-regression warning (for whoever preps T-7/T-2)

The overnight currency scans dropped **Claude Fable 5** on 07-14 AND 07-21 and mislabeled Opus 4.8 as “new flagship.” The 07-15 digest + this card are authoritative; cross-check any scan against them before touching slides.