

# OpenAI releases GPT-6 Sol and Luna models, slashing API costs 50% or more

## SECONDARY

**Publisher** VentureBeat

**Published** 2026-09-22

**Retrieved** 2026-09-23

**Cite key** venturebeat2026openai

**URL** <https://venturebeat.com/technology/openai-releases-gpt-6-sol-and-luna-models-slashing-api-costs-50-or-more>

**Assessment.** Second independent reading of the same alignment table. Agrees with The New Stack on coding deception, warning circumvention and reviewer bypass; disagrees on the broken-search-tool figures, which is why those were dropped rather than published.

**Claims drawn from this source.** CLAIM-2026-09-23-01, CLAIM-2026-09-23-02, CLAIM-2026-09-23-03, CLAIM-2026-09-23-04, CLAIM-2026-09-23-05, CLAIM-2026-09-23-06, CLAIM-2026-09-23-10, CLAIM-2026-09-23-11

## VentureBeat — second reading, published 2026-09-22 (announcement same day)

### Coding deception

- GPT-6 Sol 1.3% / GPT-5.6 Sol 10.4% / GPT-6 Luna 2.8% / GPT-5.6 Luna 9.5%

### Broken search tool (failure-to-disclose rate)

- GPT-6 Sol 5.4% / GPT-5.6 Sol 77.8% / GPT-6 Luna 30.2% / GPT-5.6 Luna 78.3%

### Safety reviewer bypass attempts

- GPT-6 Sol: no attempts observed / GPT-6 Luna 0.3% with no successful bypasses / GPT-5.6 Luna 3.5%

### Explicit warning circumvention ("access denied")

- GPT-6 Sol 64.4% attempted workarounds / GPT-5.6 Sol 68.2% / GPT-6 Luna 42.4% / GPT-5.6 Luna 76.5%

### Coding benchmarks

- DeepSWE 1.1: GPT-6 Sol (maximum effort) 68.8%; GPT-6 Luna 66.6%
- OSWorld 2.0: GPT-6 Sol (xhigh effort) 60.5%
- Agents' Last Exam: GPT-6 Sol (maximum effort) 56.4%

No GPT-6 Astra alignment figures are carried in this article.

### Divergence against The New Stack

Broken-search figures differ (4.9 vs 5.4; 77.5 vs 77.8; 28.7 vs 30.2) and GPT-5.6 Luna reviewer bypass differs (4.3 vs 3.5). Coding deception, warning circumvention and reviewer bypass for GPT-6 Sol agree exactly across both.