

OpenAI releases GPT-6.1 Sol, near GPT-6 Astra on coding at a fifth of the price

Source copy kept by wheresthe.ai for the story "OpenAI releases GPT-6.1 Sol, near GPT-6 Astra on coding at a fifth of the price" (<https://wheresthe.ai/d/openai-gpt-6-1-sol-release-2026-09/>).

Source: <https://openai.com/index/introducing-gpt-6-1-sol/>

Text saved 18 Oct 2026 (the page text through a reader, because the site refused a direct fetch).

SHA-256 of the saved text: 20a38174161c0fc619f584f608f1928745924506970add98dbab73f26c7957ae

We're introducing **GPT-6.1 Sol**, an upgrade to GPT-6 Sol that nearly matches GPT-6 Astra's intelligence on agentic coding, computer use, and professional work at one-fifth of Astra's standard input and output token prices. Cached input costs just **\$0.10 per million tokens**—95% less than standard input pricing and 50% less than GPT-6 Sol's cached input pricing—giving developers more room to build and run capable agents that reuse context across requests.

GPT-6.1 Sol offers a new balance of capability and cost for important everyday work. It delivers substantial improvements over GPT-6 Sol across complex professional tasks, from writing and debugging code to understanding documents and executing multi-step business workflows. On several of these evaluations, it approaches GPT-6 Astra's performance at substantially lower cost.

On **DeepSWE v1.1**, which evaluates complex software-engineering tasks in real codebases, GPT-6.1 Sol matches GPT-6 Astra at roughly one-fifth of the cost, while eclipsing GPT-6 Sol's best score by 6.4 percentage points at a lower reasoning effort and cost.

On **GDP.pdf**, which measures how accurately models answer professional questions using complex PDF documents, including tables, charts, diagrams, and fine-print details, GPT-6.1 Sol scores higher than Opus 5.5 with fallbacks at less than half the cost per task across the tested reasoning settings. It also approaches GPT-6 Astra's state-of-the-art performance at roughly one-fifth the cost per task.

On **AutomationBench**, which measures whether agents correctly complete multi-step business workflows, GPT-6.1 Sol scores 2.2 percentage points above Opus 5.5 at medium reasoning effort, at roughly a third of the cost. That score is also up 4.8 percentage points from GPT-6 Sol at the same setting.

GPT-6.1 Sol also makes substantial progress on tasks that require interacting with computer applications. On **OSWorld 2.0**'s offline set, which evaluates agents on demanding computer-use workflows, GPT-6.1 Sol outperforms GPT-6 Sol by seven percentage points at maximum reasoning effort at less than half the cost. It comes within 2.1 percentage points of Astra's score at maximum reasoning effort at roughly one-seventh the cost per task.

On **Terminal-Bench Science 0.1**, which evaluates scientific workflows including data analysis, simulation, and theorem proving, GPT-6.1 Sol more than doubles GPT-6 Sol's score at maximum reasoning effort at less than half the cost per task. At maximum effort, GPT-6.1 Sol costs \$5.47 per task on average, compared with \$23.21 for Opus 5.5 and \$23.80 for Astra, delivering substantial scientific capability at over 75% lower cost than either model.

GPT-6 Astra still achieves the highest score among the models tested at 68.1%, and should be used for the most difficult scientific research tasks.

GPT-6.1 Sol also improves factual accuracy on difficult prompts. Its largest factuality improvement over GPT-6 Sol comes at low reasoning effort, where it reduces the share of responses containing a factual error from 11.4% to 7.7%—a reduction of approximately 32%. Across the tested reasoning settings, its error rate remains within 1.9 percentage points of GPT-6 Astra's, at less than one-fifth the cost per task.

This evaluation measures the share of answers containing at least one factual error on de-identified conversations where users flagged an earlier model's error. These deliberately difficult prompts are not representative of typical usage.

GPT-6.1 Sol shows substantial improvements over GPT-6 Sol in our alignment evaluations, bringing it closer to GPT-6 Astra.

GPT-6.1 Sol is more transparent about its limitations and more reliable at respecting user intent and safety constraints. In challenging evaluations, it shows lower failure rates than GPT-6 Sol on transparency about broken search tools, respecting explicit restrictions, and avoiding unauthorized outcomes during agentic tasks. We observed no attempts to bypass an automated safety reviewer, matching GPT-6 Astra and GPT-6 Sol. Full details can be found in the GPT-6.1 Sol system card addendum([opens in a new window](#)).

The evaluations below deliberately test challenging situations and do not measure failure rates in typical use.

Pricing and availability

GPT-6.1 Sol is available starting today to all Plus, Pro, Business, Enterprise, and Edu users in ChatGPT Work and Codex. GPT-6.1 Sol is not yet available in Chat. Developers can also access it through the OpenAI API as `gpt-6.1-sol`. Its standard API prices are \$2 per million input tokens, \$0.10 per million cached input tokens, and \$10 per million output tokens. In the coming days, we'll also offer GPT-6.1 Sol Ultrafast, with up to 8x faster token generation compared to its standard speed in Codex.