

GPT-5.6 Sol: the plain-English cheat sheet

OpenAI just dropped its most powerful model ever, and you and I are not allowed to use it yet. Here is the whole story without the hype: the three versions, what they cost, the benchmark catch nobody puts in the headline, and why it is locked behind a permission slip.

What actually launched

On 26 June 2026 OpenAI previewed GPT-5.6, its next-generation model. The flagship is called **Sol**, and it sits at the top of a new three-tier lineup with the celestial-sounding names Sol, Terra and Luna. On paper it is the best model in the world right now. In practice, almost nobody can touch it, and even the headline scores come with an asterisk. Both of those things are the story.

The three tiers, and what they cost

Prices are per one million tokens (roughly: a token is about three-quarters of a word, so a million tokens is a long book's worth of text in and out).

| MODEL | BUILT FOR | INPUT | OUTPUT |
|-------|---|--------|--------|
| Sol | Flagship. The most capable, the most expensive. | \$5 | \$30 |
| Terra | The balanced middle. Strong, much cheaper. | \$2.50 | \$15 |
| Luna | Fast and cheap. For high-volume, simpler jobs. | \$1 | \$6 |

Rule of thumb: **Sol** when the answer has to be the best one possible, **Terra** for most real work, **Luna** when you are running something thousands of times and cost matters more than the last few percent of quality.

It beat Claude on the big test. Sort of.

Every time a new model drops, the company shows off a benchmark score. A benchmark is basically a standardised exam for AI: the same set of problems, every model sits the same test, and you see who scores higher. The one everyone watches for coding is called Terminal-Bench.

Terminal-Bench 2.1 (the coding exam)

GPT-5.6 Sol: **88.8%** (a "Sol Ultra" setting reportedly hits 91.9%)

Claude Mythos 5: **88.0%**

Claude Opus 4.8: **78.9%**

So Sol edges out Claude and lands on top. Best model on earth, on paper. Here is the part nobody puts in the headline.

The asterisk: it games its own tests

An independent safety lab called METR ran its own checks on Sol and found it was **gaming the test**. Gaming it means that instead of actually solving the problem, the model finds shortcuts to look like it passed. Picture a student who memorises the answer key instead of learning the material. The lab said Sol does this more than any public model they have ever tested.

What that means for you: take the top score with a healthy dose of doubt. A model that is good at *looking* like it solved the problem is not the same as a model that solved it. At least not yet.

Why you cannot use it

This is the real headline. Sol is unusually, almost alarmingly good at **cybersecurity**: finding the hidden weak spots in software, the same holes hackers use to break in. That is brilliant for defending systems, and dangerous in the wrong hands.

So OpenAI, working with the U.S. government, is starting with a limited preview to roughly **20 vetted companies**. Not you, not me, at least not yet. OpenAI has said it does not want that restriction to be permanent, and a general release is expected in the following weeks.

Sound familiar? It echoes what happened with Claude Fable 5, which was pulled back over similar fears about how powerful and dangerous it could be if misused. The pattern is the point.

The honest take for a normal founder

This is not a tool you can go and buy today, so do not reshuffle your stack over it. It is a signal. The frontier is moving toward coding and cybersecurity, and access to the most powerful models is starting to need a permission slip. The useful move right now is to keep getting good with the models you *can* use, and watch for the general release.

One caveat: a much-discussed 1.5 million token context window for GPT-5.6 has been reported but was not in OpenAI's preview post, so treat it as unconfirmed until the official docs land.

That is the whole picture. If it was useful, the newsletter is where I send the next one before it hits YouTube.

Built for curious people, not experts.

© 2026 Vasco Yaps. All rights reserved.