

Grok 4.6 Matches GPT-5.6 Sol, and Grok Bot Gives Agents Their Own Cloud Computer

Kabui, Charles

2026-08-21

[Read at ToKnow.ai](#)



SpaceXAI released Grok 4.6 and Grok Bot on August 11 and 12, two pieces of the same bet: AI that works like a team, not a chat window. Grok 4.6 scores 61 on the [Artificial Analysis Intelligence Index](#), a composite of nine benchmarks, matching GPT-5.6 Sol and trailing only Claude Fable 5's 62. It also beats Sol on CursorBench 3.2, a coding-agent benchmark, 69.9%

to 67.2%. It costs \$2 per million input tokens and \$6 per million output tokens, with a faster variant at double the price.

Each Bot runs on a persistent cloud computer of its own, a virtual machine with a browser, filesystem, and terminal, so it can sign into apps you already use, even ones with no API, and finish jobs end to end. It keeps working after you close your laptop and only returns when something needs approval. Show it a workflow once and it re-runs it as a routine. The beta is available for SuperGrok Heavy, Cursor Ultra, and Cursor Teams Premium on desktop and iOS. One caveat: all your Bots share a single cloud computer per account, and SpaceXAI's docs warn against treating separate Bots as a security boundary.

Agents become first-class infrastructure with their own compute, and the no-API problem is solved by driving the UI the way a person would. SpaceX closed its [Cursor deal](#) the same week, so it now owns the model, the IDE, and the computers agents live on.

Read More: [Grok 4.5: Fast and Cheap, but Not the Best at Every Coding Test](#)

Sources:

- [Introducing Grok 4.6 \(SpaceXAI\)](#)
- [Introducing Grok Bot \(SpaceXAI\)](#)
- [Grok Bot overview \(SpaceXAI Docs\)](#)
- [Grok Bot: Approvals, security, and privacy \(SpaceXAI Docs\)](#)
- [Grok 4.6 model analysis \(Artificial Analysis\)](#)

***Disclaimer:** For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)*