

xAI releases Grok 4.7 on a larger base model at Grok 4.6 prices

Source copy kept by wheresthe.ai for the story “xAI releases Grok 4.7 on a larger base model at Grok 4.6 prices” (<https://wheresthe.ai/d/xai-grok-4-7-release-2026-09/>).

Source: <https://x.ai/news/grok-4-7>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: c6f1bffe22cc422f3daa1867b919785ed18499d0d61ac30a202b92c3b4923d0

Introducing Grok 4.7 | SpaceXSI Back to news Sep 21, 2026 Introducing Grok 4.7 SpaceXSI's most powerful model for coding and knowledge work. Twice as fast, at half the price of comparable models. Try for free Start building

Grok 4.7 is our most capable model for coding and knowledge work. It works longer on difficult tasks, checks its own work more carefully, and comes with our best-calibrated safeguards to date. Served at the same price and speed as Grok 4.6 , it is highly competitive in its class. Cost Tokens Steps

On CursorBench 4.0, which stresses longer-running coding tasks, Grok 4.7 is at the frontier in price-performance. Model Improvements Grok 4.7 uses a new, larger base model compared to Grok 4.6 . It was trained with a longer reinforcement learning run on a harder mix of tasks, weighted toward problems that take many hours to complete. The model is better at verifying its own work and managing longer context. We also trained Grok 4.7 to natively understand the Grok Bot harness, making it better at conversational tasks and general knowledge work.

Grok 4.7 xHigh Grok 4.6 High GPT-5.6 Sol Max Fable 5.1 Max

Input token price \$ per million

\$2 \$2 \$4 \$10

Output token price \$ per million

\$6 \$6 \$20 \$50

Software engineering CursorBench 4.0

46.3% 40.4% 41.7%

51.8%

Software engineering DeepSWE v1.1

71.0%* 65.2%

72.7% 70.0%

Electrical engineering EEBench

64.0% 53.0% 39.4% 56.4%

Multi-hour office work AA Briefcase v1.1

1,657 1,546 1,487

1,678

Multi-hour terminal work Terminal-Bench 4.0

37.6% 20.3% 37.3%

57.9%

Legal work Harvey Legal Agent Benchmark

19.6% 15.8% 2.5% 6.7%

Clinical reasoning HealthBench Professional

56.7% 48.5% 60.5%

62.1%

* high effort Token prices and benchmark scores for Grok 4.7, Grok 4.6, GPT-5.6 Sol, and Fable 5.1. Benchmarks are CursorBench 4.0, DeepSWE v1.1, EEBench, AA Briefcase v1.1, Terminal-Bench 4.0, Harvey Legal Agent Benchmark, and HealthBench Professional. An asterisk on Grok 4.7 DeepSWE marks a high-effort score. Grok 4.7 is better at creating documents and presentations. In GDPval and AA Briefcase, AI is asked to work on tasks done by professionals such as lawyers, nurses, and financial analysts. Grok 4.7 improves upon Grok 4.6 on both benchmarks and performs comparably to other frontier models. Professional knowledge work Multi-hour office work Electrical engineering

Professional knowledge work GDPval

GDPval, AA Briefcase, and EEBench scores comparing Grok 4.7 with Grok 4.6, Fable 5.1, and GPT-6 Astra. Safety & Cybersecurity Grok 4.7 was built with an entirely new safeguard stack. It is the strongest model we've tested on refusals and jailbreak resistance. In dual-use domains like cybersecurity and biological work, it leads on both utility for benign tasks and safe refusal on dangerous ones, topping LatchBio's biosafety benchmark at 62.4%. Grok 4.7 balances strong cyber defense capabilities with low refusal rates for legitimate use. It shows the highest safety on HackerBench v0.3, our benchmark for risky and malicious cyber tasks, allowing only 3.3% of risky dual-use prompts through while rarely blocking legitimate security work. We've also started giving select cybersecurity partners invite-only access to Grok 4.7's red-team capabilities for defense research. Pricing and availability Grok 4.7 is available today in Cursor and Grok Build . It is also available through the Grok API , third-party coding harnesses, and model routers and cloud platforms. The model is priced starting at \$2 per million input tokens and \$6 per million output tokens. We also serve a fast variant with twice the output speed at twice the price.

Console Create an API key

docs.x.ai Read the docs

Try it in Grok Build for free Get started today at x.ai/build . \$ curl -fsSL https://x.ai/cli/install.sh | bash