

xAI releases Grok 4.5, trained alongside Cursor

Source copy kept by wheresthe.ai for the story "xAI releases Grok 4.5, trained alongside Cursor" (<https://wheresthe.ai/d/xai-grok-4-5-release-2026-07/>).

Source: <https://x.ai/news/grok-4-5>

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

SHA-256 of the saved text: 0408b3c4f2eb1a319dcd464f95b3e5c85c24b0d5ab7c029d125c0f5f6e7584fa

Introducing Grok 4.5 | SpaceXSI Back to news Jul 16, 2026 Introducing Grok 4.5 Grok 4.5 is SpaceXSI's smartest model built for coding, agentic tasks, and knowledge work. Try for free Start building

Today, we're launching Grok 4.5 , SpaceXSI's smartest model built to excel at coding, agentic tasks, and knowledge work. It's our strongest model ever and was trained alongside Cursor . Real-world engineering excellence Grok 4.5 was trained on datasets spanning knowledge in coding, science, engineering, and math. With both intelligent and efficient reasoning, Grok 4.5 excels at real engineering tasks and exceeds comparable leading models at these tasks. DeepSWE 1.0 DeepSWE 1.1 SWE Marathon Terminal Bench 2.1 SWE Bench Pro

Competitor figures are drawn from the respective developers' published system cards or benchmark leaderboards

Benchmark bar charts of model scores. DeepSWE 1.0 (eval created by Datacurve, run with each model provider's harnesses by AA): Fable (max) 66.1%, GPT 5.5 (xhigh) 64.31%, Grok 4.5 62.0%, Opus 4.8 (max) 55.75%, Opus 4.7 (max) 40.12%. DeepSWE 1.1 (mini-swe-agent harness run by Datacurve): Fable (max) 70%, GPT 5.5 (xhigh) 67%, Opus 4.8 (max) 59%, Grok 4.5 53%, GLM 5.2 44%. SWE Marathon resolution rate (pass@1): Grok 4.5 29.0%, Opus 4.8 (max) 26.0%, Fable (max) 24.0%, Opus 4.7 (max) 16.0%. Terminal Bench 2.1: Fable (max) 84.3%, GPT 5.5 (xhigh) 83.4%, Grok 4.5 83.3%, Opus 4.8 (max) 78.9%, Opus 4.7 (max) 78.9%. SWE Bench Pro resolve rate: Fable (max) 80.4%, Opus 4.8 (max) 69.2%, Grok 4.5 64.7%, Opus 4.7 (max) 64.3%, GLM 5.2 62.1%, GPT 5.5 (xhigh) 58.6%. Competitor figures are drawn from the respective developers' published system cards or benchmark leaderboards. Training Grok 4.5 Grok 4.5 was trained across tens of thousands of NVIDIA GB300 GPUs, with training and stability techniques designed for large-scale runs. Beyond raw token volume, we invested heavily in data filtering and curation: deduplication, quality scoring, and domain-focused selection so that the data mixture stayed high-coverage and high-signal. We scaled reinforcement learning with a strong focus on per-token intelligence. Our RL training covers hundreds of thousands of tasks, centered on multi-step software engineering and other technical work, with automated and model-based grading. Our stack is built for highly asynchronous training, so agentic rollouts can run for many hours while learning continues across tens of thousands of GPUs. The result is more intelligent and efficient reasoning on real engineering and agentic tasks. Built with one prompt Grok 4.5 is incredibly capable at coding, from challenging Rust and C/C++ tasks to end-to-end app building from prompt to production. Below are some examples built by the model with one prompt. Grok 4.5 is highly proficient at creating well-designed, end-to-end functional apps even with minimal specification. Solar system

Make a beautiful simulation of the universe and solar system. should be sped up with adjustable time, realistic motion, orbits, stars. use threejs. Make the HUD well styled and conform to modern design principles.

Make a beautiful simulation of the universe and solar system. should be sped up with adjustable time, realistic motion, orbits, stars. use threejs. Make the HUD well styled and conform to modern design principles.

app.localhost — cosmos

Loading solar system...

Faster than flash models Grok 4.5 is served at fast-model speeds of 80 TPS . Combined with twice greater token efficiency than the latest leading models at the same tasks, the model delivers intelligent results to you more quickly and at far lower costs. Token efficiency avg. output tokens per SWE Bench Pro task

Grok 4.5 :host{display:inline-block;direction:ltr;white-space:nowrap;line-height:var(--number-flow-char-height, 1em) !important}span{display:inline-block}:host([data-will-change]) span{will-change:transform}.number,.digit{padding:calc(var(--number-flow-mask-height, 0.25em) / 2) 0}.symbol{white-space:pre} 0 0

Opus 4.8 (max) :host{display:inline-block;direction:ltr;white-space:nowrap;line-height:var(--number-flow-char-height, 1em) !important}span{display:inline-block}:host([data-will-change]) span{will-change:transform}.number,.digit{padding:calc(var(--number-flow-mask-height, 0.25em) / 2) 0}.symbol{white-space:pre} 0 0

4.2× fewer tokens

0 70 k tokens

Token efficiency, average output tokens per SWE Bench Pro task — Grok 4.5 resolves tasks with 15,954 output tokens on average, about 4.2× fewer than Opus 4.8 (max) at 67,020. Excels at Office work Grok 4.5 is now the default model in Grok Build . In addition to its coding proficiency, Grok Build is capable of building complex Excel models that involve research from the web, multi-sheet formula use, and even leaves stickies or notes behind for future reference. In PowerPoint and Word, Grok 4.5 is similarly meticulous. The model is capable of using native PowerPoint shapes to build complex diagrams, designing intuitive slide content, and writing clear prose in Word. Outline a 5-slide quarterly business review

AutoSave Q3 Review.pptx Search

Home Insert Design Transitions Animations Review Comments Share

New Slide Layout B I U Arrange Design Ideas Add-ins Grok 1

Quarterly review Oct · FY26 Q3 Business Review Revenue, margins, pipeline — and where we invest next

01 Revenue 02 Margins 03 Pipeline 04 Outlook

Slide 1 of 5 · English (US) 100%

Learn more about our plugins for Word , PowerPoint , Excel , and Outlook . Pricing Grok 4.5 is delivered at an incredibly competitive cost compared to other leading models. Grok 4.5 is priced at \$2 per million input tokens and \$6 per million output tokens. The model also achieves roughly 2x the token efficiency of comparable leading models, solving tasks in under half the number of steps. Overall, Grok 4.5 delivers the highest intelligence per unit of time and cost. Getting started Grok 4.5 is available today in Grok Build, in Cursor on all plans, and from the SpaceXSI console . Simply grab an API key and get started in a few lines of code:

```
Copy curl -s https://api.x.ai/v1/responses \ -H "Authorization: Bearer $XAI_API_KEY" \ -H "Content-Type: application/json" \ -d '{ "model": "grok-4.5", "input": "Find and fix the bug, then explain it: function median(a) {a.sort();return a[a.length/2]}"}'
```

bash

Console Create an API key

docs.x.ai Read the docs

Try it in Grok Build for free We're offering free Grok 4.5 usage for a limited time in Grok Build and Cursor. Get started today at x.ai/build . `$ curl -fsSL https://x.ai/cli/install.sh | bash`