

Grok 4.6

xAI · detected 2026-08-12 15:54 UTC · alert out 63s later · 106 alerts delivered

AT A GLANCE

PRICE PER 1M

\$2 / \$6

input / output · cached
\$0.50

CONTEXT

500K

tokens in

MAX OUTPUT

450K

tokens per reply

TAKES


text, image, file

returns text


HOW IT COMPARES

MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
Grok 4.6	new	\$2	\$6	500K	60
Grok 4.7	it follows	\$2	\$6	500K	67.1
Claude Opus 5.5	today's #1	\$4	\$20	1M	100
Muse Spark 1.3	best value	\$1.25	\$4.25	1.05M	79.1
vs Grok 4.7	same price the model it follows: \$2 / \$6, bench index 67.1				
vs Claude Opus 5.5	input 50% cheaper, output 70% cheaper today's #1: \$4 / \$20, bench index 100; context 1M				
vs Muse Spark 1.3	input 60% more, output 41% more best value: \$1.25 / \$4.25, bench index 79.1; context 1.05M				

PRICE PER 1M TOKENS, BLENDED 3 IN : 1 OUT



WHERE THE PUBLIC LEADERBOARDS STAND TODAY



Grok 4.6 sits at 60 on the bench index. The bench index is the median of each model's percentile across the public leaderboards that rate it (LMArena, LiveBench, SimpleBench, LLM Stats, Vals AI, Artificial Analysis).

Sources: price, limits and modalities for Grok 4.6 as the OpenRouter catalog lists x-ai/grok-4.6; the other models as their makers list them (models.dev); rankings from the ArtificialWatch Bench Index. Facts only: no forecast of ours, no claims from the lab's announcement.