Command A Plus

Cohere

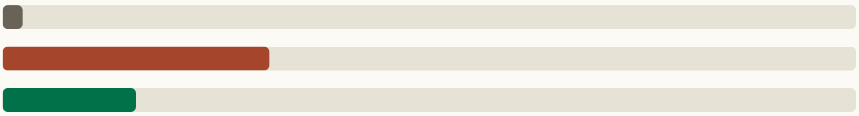
AT A GLANCE

PRICE PER 1M	CONTEXT	MAX OUTPUT	TAKES
$0.30 / $1.50	**192K**	**64K**	**text, image**
input / output · cached $0.15	tokens in	tokens per reply	returns text

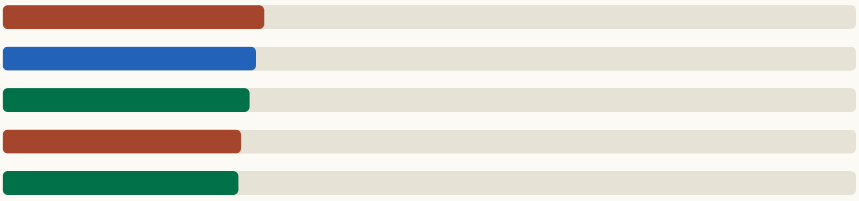
HOW IT COMPARES

MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
Command A Plus	new	$0.30	$1.50	192K	not rated yet
Claude Opus 5.5	today's #1	$4	$20	1M	98.1
GPT-6.1 Sol	best value	$2	$10	1.05M	88.4
vs Claude Opus 5.5	**93% cheaper**	today's #1: $4 / $20, bench index 98.1; context 1M			
vs GPT-6.1 Sol	**85% cheaper**	best value: $2 / $10, bench index 88.4; context 1.05M			

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



WHERE THE PUBLIC LEADERBOARDS STAND TODAY



Command A Plus is not rated yet: the public boards usually score a new model within days, and its page on artificialwatch.com gains the scores then. 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 Command A Plus as the OpenRouter catalog lists cohere/command-a-plus; 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.