

Gemini 3.6 Flash

Google

AT A GLANCE

PRICE PER 1M

\$0.75 / \$3.75

input / output · cached
\$0.075

CONTEXT

1.05M

tokens in

MAX OUTPUT

66K

tokens per reply

TAKES


**text, image, video,
file, audio**

returns text


HOW IT COMPARES

MODEL	ROLE	IN / 1M	OUT / 1M	CONTEXT	BENCH INDEX
Gemini 3.6 Flash	new	\$0.75	\$3.75	1.05M	not rated yet
Gemini 3.8 Flash	it follows	\$0.75	\$3.75	1.05M	44
Claude Opus 5.5	today's #1	\$4	\$20	1M	100
Claude Sonnet 5.5	best value	\$2	\$10	1M	89.8
vs Gemini 3.8 Flash	same price	the model it follows: \$0.75 / \$3.75, bench index 44			
vs Claude Opus 5.5	81% cheaper	today's #1: \$4 / \$20, bench index 100; context 1M			
vs Claude Sonnet 5.5	63% cheaper	best value: \$2 / \$10, bench index 89.8; context 1M			

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



WHERE THE PUBLIC LEADERBOARDS STAND TODAY



Gemini 3.6 Flash 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 Gemini 3.6 Flash as the OpenRouter catalog lists google/gemini-3.6-flash; 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.