

Which AI Model Fits Which Budget? Intelligence per Dollar in 2026

MANAGEMENT SUMMARY & CORE TAKEAWAY

AI procurement is shifting in 2026 from a performance race to a cost-efficiency calculation. Prices between models span a factor of 1,250 – from \$0.06 to \$75 per million tokens. What matters isn't the most powerful model, but the one with the best price-performance ratio for your specific use case: customer support, document analysis, or code generation.

CONTEXT & EXECUTIVE RELEVANCE

For Swiss SMEs in 2026, the question is no longer which model tops the leaderboard, but how much capability a franc actually buys. The price gap between the cheapest and most expensive top-tier model runs to 1,250 times – from \$0.06 to \$75 per million tokens. Companies that allocate budget by use case rather than by hype can cut operating costs for typical SME workloads such as customer support or document analysis by up to 90 percent, without a noticeable drop in quality.

A year ago, the debate centred on which model was 'smartest'. A widely shared LinkedIn chart from late July 2026 captures the shift precisely: the metric that matters is no longer 'who is smartest' but 'intelligence per dollar'. GPT-5.5 and GPT-5.6, along with Gemini, position themselves as price-performance leaders in that chart, while the Chinese model Kimi emerges as a serious challenger. For decision-makers, this means model selection is no longer a ranking contest – it is a procurement question.

CORE ACTION ITEMS

- 01 GPT-4.1 Nano: \$0.10 / \$0.40 per million tokens (input/output)
- 02 Gemini Flash: \$0.10 / \$0.40 per million tokens
- 03 Llama 4 Scout: \$0.08 / \$0.18 – the cheapest open-weight model (AI Security Gateway, April 2026)
- 04 Qwen 3: \$0.06 / \$0.12 – a budget option for multilingual applications

SWISS COMPLIANCE & ARCHITECTURE GUIDELINES

- ☑ Full compliance with revised Swiss Data Protection Act (revDSG).
- ☑ Zero vendor lock-in on closed proprietary models.
- ☑ Seamless integration into Abacus, Bexio and M365 ecosystems.

Non-Binding Process Analysis

Let us review in 20 minutes where AI agents create immediate measurable efficiency in your workflows.

ki-outsourcing.ch