

Kimi K3 vs. GLM-5.2: Capability, Price, Openness, and Deployment

Kabui, Charles

2026-07-20

Table of Contents

Introduction	2
1. Capability favours Kimi	2
2. Price and access favour GLM	3
3. “Open” is the decisive difference	3
4. Conclusion	3
References	4

 Read at ToKnow.ai



Introduction

Kimi K3 is the stronger model in today's independent comparison. GLM-5.2 is the easier model to buy, host, and control. Both accept about one million tokens, but only K3 reads images and video. GLM-5.2 takes text only and allows up to 128,000 output tokens ([Kimi quickstart](#), [GLM-5.2 guide](#)).

1. Capability favours Kimi

Artificial Analysis scores K3 at **57 versus 51** on its Intelligence Index. OpenRouter's live comparison also shows K3 ahead in coding, **76 versus 69**, and agent tasks, **50 versus 43** ([Kimi evaluation](#), [GLM evaluation](#), [live comparison](#)).

These are useful, not final. The Intelligence Index is an English, text-only suite weighted toward agents, coding, and scientific reasoning. It cannot measure K3's visual advantage, and a six-point lead may not predict results in your repository ([methodology](#)). Vendor tables are less comparable because Kimi and Z.ai sometimes use different agent software, reasoning settings, and run counts ([Kimi technical blog](#), [GLM-5.2 announcement](#)).

2. Price and access favour GLM

Checked July 20, 2026	Kimi K3	GLM-5.2
Official input, per million tokens	\$3.00	\$1.40
Official cached input	\$0.30	\$0.26
Official output	\$15.00	\$4.40
OpenRouter providers (live)	1	28

GLM output is about **71% cheaper** at official list prices ([Kimi pricing](#), [Z.ai pricing](#)). OpenRouter currently lists an even cheaper GLM route, but provider prices, latency, and throughput move constantly. Treat that page as a live quote, not a permanent specification.

3. “Open” is the decisive difference

GLM-5.2 has downloadable **753-billion-parameter** weights under the MIT licence, plus documented support for vLLM, SGLang, Transformers, KTransformers, and other runtimes ([model card](#)). That does not make it a laptop model. The full checkpoint and its one-million-token cache still demand serious infrastructure.

K3 has 2.8 trillion total parameters and activates 16 of 896 experts for each token. Moonshot recommends at least 64 accelerators for deployment. More importantly, its full weights are only promised for July 27. On July 20, K3 is an API model, not an open-weight release, regardless of the vendor’s “open” wording ([technical blog](#)).

4. Conclusion

Choose K3 for the best current scores, native vision, and difficult agent work. Choose GLM-5.2 for lower bills, provider choice, downloadable weights, and deployment control. Test both in the same harness before committing because the agent software around the model can change the result.

- **Capability:** Kimi K3
- **Cost and availability:** GLM-5.2
- **Ownership:** GLM-5.2, unless K3’s promised release arrives with usable terms

For more background, read our [GLM-5.2 model profile](#).

References

1. Moonshot AI (2026). *Kimi K3: Open Frontier Intelligence*.
2. Moonshot AI (2026). *Kimi K3 API quickstart*.
3. Moonshot AI (2026). *Kimi K3 pricing*.
4. Z.ai (2026). *GLM-5.2: Built for Long-Horizon Tasks*.
5. Z.ai (2026). *GLM-5.2 model guide*.
6. Z.ai (2026). *Model pricing*.
7. Z.ai (2026). *GLM-5.2 model card and weights*.
8. Artificial Analysis (2026). *Kimi K3 evaluation*.
9. Artificial Analysis (2026). *GLM-5.2 evaluation*.
10. Artificial Analysis (2026). *Intelligence benchmarking methodology*.
11. OpenRouter (2026). *Kimi K3 vs. GLM-5.2 live comparison*.

Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. **Read more:** [/terms-of-service](#)