


The AI race: Why “Mythos” is not the end of the story

By Frank , 17 April 2026

The recent announcement of Anthropic’s new AI model, Mythos, has caused quite a stir. With claims that it can autonomously detect thousands of cybersecurity vulnerabilities — including in every major operating system and every browser — the model is being praised as a game changer.



©Pixabay

Even U.S. Treasury Secretary Scott Bessent called it a “revolutionary step” that would give the United States a lead of 3 to 6 months over China in the global AI race.

But is the story really that simple?

What is Mythos — and what is it not?



Mythos is an early preview of a general AI model with strong reasoning and coding capabilities. It is being deployed specifically within Project Glasswing, a collaboration between Anthropic and twelve partners such as Apple, Amazon, Microsoft, Nvidia, and JPMorgan Chase. The goal: defensive cybersecurity. Companies can use the model to identify and patch weak points in their own systems.

Important: Mythos is not publicly available.


Anthropic fears that, in the wrong hands, the model could instead become a powerful weapon for cyberattacks. So the claim that it can “hack every operating system” deserves some nuance: this concerns authorized, defensive use within strict boundaries.



The broader AI race: more than one model

Anyone who reduces the AI race to a single model misses the bigger picture. As Nvidia CEO Jensen Huang has repeatedly emphasized, China has structural advantages that should not be overlooked.

- Energy: China has a rapidly growing, low-cost energy infrastructure, including renewable sources. The country is even dealing with excess solar energy: the “wastage” percentage (unused electricity) rose from 6.1% to 9.2% in early 2026.
- For AI, where electricity accounts for 50–75% of operational costs, this is a strategic advantage.
- Construction speed: Building a data center takes an average of 3 years in the United States; in China, the same can be done in months
- Costs: Estimates suggest that building a data center in China costs around \$6 million per megawatt, compared with \$10–15 million in the U.S.
- Chips: Huawei’s new Ascend 950PR chip reportedly performs up to 2.8 times better than Nvidia’s China-approved H20 chip, at a fraction of the price. China is aiming for 80% chip self-sufficiency by 2030, an ambitious but serious goal.



Why the “five-layer cake” matters

Experts often describe the AI race as a five-layer cake.

1. Models & applications (such as ChatGPT, Mythos) — the top layer, what consumers see
2. AI factories (training & inference)
3. Chips & hardware
4. Energy infrastructure
5. Raw materials & production capacity — the foundation layer

The United States leads in the upper layers: advanced models, efficient chips, leading research. But China holds a strong position in the lower layers: energy, scale, production costs, and a vast supply of technical talent.

Conclusion: no winner in the short term

Mythos is undoubtedly a technical feat. But it is naive to think that one model will decide the AI race. As Jensen Huang of Nvidia warns: do not underestimate China. The country has the resources, the scale, and the long-term vision to catch up, including in advanced chips and models.

The real question is not: “Who wins the AI race?” But: **How do we ensure that this technology is used safely, fairly, and for the benefit of humanity?** All sides — the U.S., China, and Europe — share a responsibility in that regard.

<https://www.youtube.com/watch?v=IUHQIwrqo-8>