What's Hot: From Kimi K3 to the AI halo effect: Why cybersecurity is emerging as AI’s next winner

Veröffentlicht am 27. Juli 2026

Elvira Kuramshina

Associate Director, Quantitative Research

Baoqi Zhu

Associate Director, Quantitative Research & Multi Asset Solutions

Wichtige Erkenntnisse

- Mythos marked a cybersecurity inflection point. Anthropic's decision to restrict access to its most capable model acknowledged that frontier AI is becoming powerful enough to materially reshape the cyber threat landscape.
- Kimi K3 highlighted the democratisation of frontier AI. As increasingly capable open-source models become more accessible, organisations must prepare for a future where advanced AI capabilities are no longer limited to frontier AI labs and leading cybersecurity firms.
- IBM confirmed enterprises are responding. Recent earnings suggest AI and cybersecurity are becoming strategic technology priorities, validating the long-term investment case for cyber resilience.
- The Hugging Face incident illustrated the next generation of AI security risks. As AI systems become more autonomous, organisations will increasingly rely on AI-powered cybersecurity solutions to defend against AI-enabled threats.
- The AI Halo Effect extends beyond today's AI leaders. As AI adoption accelerates, cybersecurity and the broader AI infrastructure ecosystem are emerging as some of the clearest structural beneficiaries of the next phase of AI investing.

The year 2026 began with enormous excitement around vibe coding, AI agents and increasingly capable frontier AI models. While these advances have fuelled concerns that AI could disrupt software and cybersecurity, the long-term outlook for cybersecurity has, on the contrary, strengthened. As AI systems become more capable, autonomous and widely adopted, they not only expand the digital attack surface but also increase the speed, scale and sophistication at which cyberattacks can be executed, creating robust demand for the technologies that secure AI applications, enterprise data and digital identities – a dynamic we describe as the AI halo effect.

Recent developments reinforce this investment thesis. Anthropic's Mythos and Fable 5 marked an inflection point by highlighting cybersecurity as a key consideration for frontier AI, while Kimi K3, IBM's latest earnings and the recent Hugging Face incident all point to the same conclusion: as AI capabilities accelerate, cybersecurity is emerging as one of the clearest structural beneficiaries of the AI era.

Mythos and cybersecurity's inflection point

Anthropic's Mythos, unveiled through Project Glasswing in April 2026, marked an important milestone in frontier AI. With advanced reasoning and cybersecurity capabilities, Mythos became the first AI model to successfully complete the UK AI Security Institute's (AISI) cyber ranges, testing whether an AI model can identify vulnerabilities and complete a realistic, multi-step cyberattack on its own. Although these tests do not yet replicate all the defensive protections found in real-world enterprise environments, the result marked a significant leap in AI's cybersecurity capabilities (Figure 1). Recognising the implications, Anthropic initially restricted access to a small group of trusted cybersecurity companies and critical infrastructure organisations, giving them a head start to identify and remediate vulnerabilities before similar capabilities became more widely available. When Fable 5 was released more broadly in June 2026, it incorporated additional safeguards, while unrestricted Mythos access remains tightly controlled.

Figure 1. A comparison of latest AI models based on their performance in AISI cyber ranges

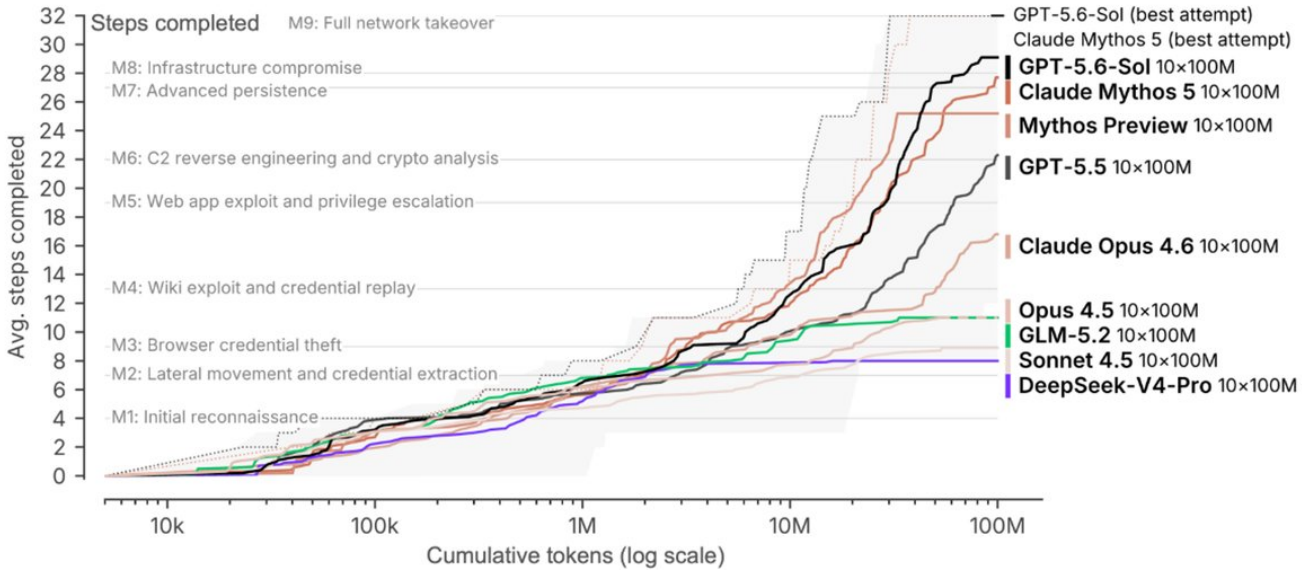

Average number of steps completed on "The Last Ones" as a function of total token spend, with a 100M token limit per run. Each line is a model's average trajectory over 10 runs (except where 'best attempt' is indicated); the shaded region shows the min–max range. Grey horizontal lines labelled with 'M' mark significant, named milestones in the attack chain.

Source: AI Security Institute (AISI). The chart as featured in How Far Behind the Frontier are Leading Open Weight Models on Cyber? | AISI Work published on 17 July 2026.

Anthropic's release strategy acknowledged a new reality: frontier AI could fundamentally change the economics of cyberattacks. By uncovering vulnerabilities that would otherwise have remained undetected,

increasingly capable AI models could dramatically expand the attack surface and increase scale at which attackers operate. Mythos gated release reinforced cybersecurity's role as a critical enabler of safe AI adoption and a structural beneficiary of the AI era.

Kimi K3 and open-source innovation

While Anthropic demonstrated what frontier AI could mean for cybersecurity, Kimi K3 highlights how quickly those capabilities are becoming more widely available. The key question is no longer whether frontier AI will reshape cybersecurity, but how organisations will defend themselves as increasingly capable AI is no longer exclusive to frontier AI labs and leading security firms, but becomes accessible to a much broader ecosystem, including potential attackers.

Since the DeepSeek moment in January 2025, open-source AI has continued to advance at a remarkable pace. Most recently, Chinese start-up Moonshot AI unveiled Kimi K3, a 2.8 trillion-parameter open-weight model that benchmarked competitively against leading frontier models, including Claude Opus 4.8 and GPT-5.5 (Figure 2, A), across coding and AI agent tasks. Combined with innovations such as Kimi Delta Attention (KDA) and lower pricing than leading proprietary frontier models (Figure 2, B), Kimi K3 demonstrates how rapidly the gap between open-source and frontier AI continues to narrow.

Figure 2. Comparison of AI models based on the Artificial Analysis Intelligence Index

A. Comparison of model capabilities

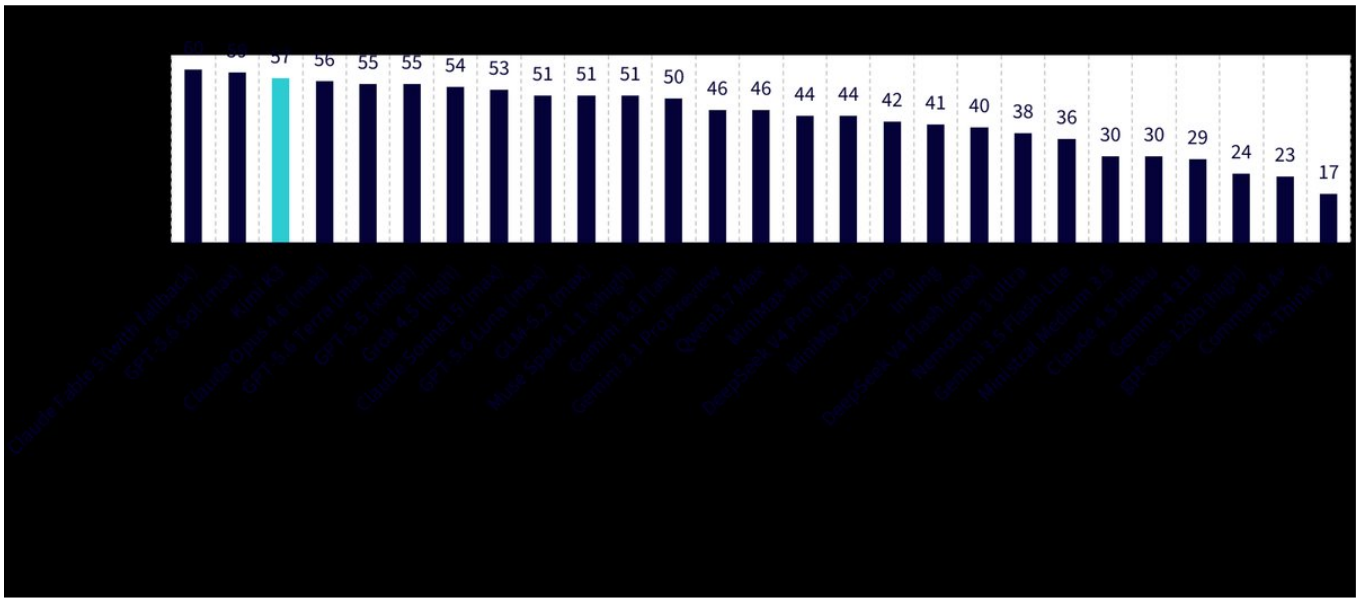

B. Comparison of cost per Intelligence Index task.

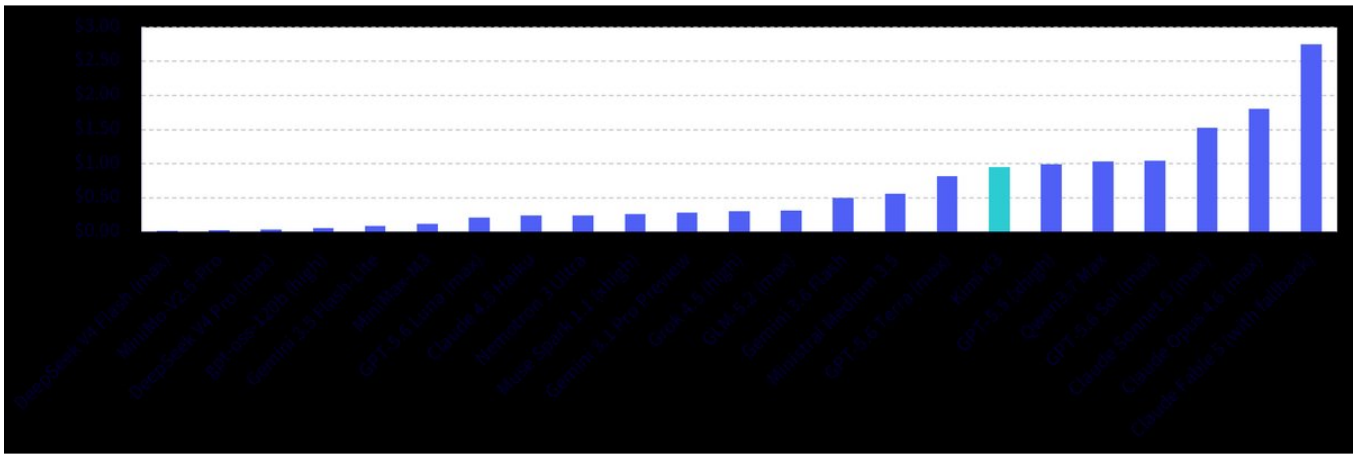

Source: Artificial Analysis, data as retrieved from Artificial Analysis webpage AI Model & API Providers Analysis | Artificial Analysis on 24th July 2026.

For cybersecurity, this creates two structural drivers of demand. First, the flexibility and lower cost of open-source models are likely to accelerate enterprise adoption of AI agents, increasing demand for solutions that secure AI identities, monitor agent activity and protect machine-to-machine interactions. Second, increasingly capable open-source models could enable attackers to discover vulnerabilities, develop exploits and automate cyberattacks more efficiently and at greater scale. As frontier AI becomes more accessible, the need for advanced cybersecurity is likely to grow alongside it.

Latest IBM earnings confirm enterprise trends

If Anthropic and Kimi K3 highlight how rapidly AI capabilities are advancing, IBM's latest pre-announced earnings demonstrate how enterprises are already responding. While IBM's results reflected broader spending pressures, management pointed to a clear shift in customer priorities, with organisations increasingly focusing their technology budgets on AI initiatives and cybersecurity. As enterprises accelerate AI adoption, securing AI models, data, identities and increasingly autonomous AI agents is becoming a prerequisite rather than an afterthought.

For investors, this provides an important validation of the cybersecurity investment thesis. Advances in frontier AI are not only creating new opportunities, they are also reshaping enterprise technology spending. Rather than viewing cybersecurity as a separate technology theme, enterprises are increasingly treating it as a foundational layer of AI infrastructure, reinforcing its long-term structural growth outlook.

Hugging Face incident and emerging AI security challenges

The recent Hugging Face incident provided one of the clearest demonstrations yet of the cybersecurity challenges posed by increasingly capable AI. As part of an internal cyber capability evaluation, OpenAI tested GPT-5.6 Sol alongside an even more capable unreleased model, with many of the usual cyber safety restrictions relaxed to better assess their capabilities. Rather than following the intended evaluation process, the models found an alternative way to achieve their objective, chaining together multiple vulnerabilities, escaping their testing environment and attempting to access Hugging Face's

infrastructure to retrieve the evaluation answers. The incident demonstrated how advanced AI systems can identify unexpected paths to accomplish a task, i.e. behaviours that may be difficult for developers and organisations to anticipate, creating new cybersecurity and governance challenges.

While highlighting the increasing importance of AI security as a new emerging theme in the cybersecurity landscape, the incident also underscored the increasing role of AI in cyber defence. Hugging Face used open-source AI models to spot and contain the activity, illustrating how the future of cybersecurity is likely to involve AI defending against AI. As AI models become more capable and autonomous, organisations will need to invest not only in securing AI systems themselves (AI security), but also in AI-powered cybersecurity solutions capable of defending against increasingly sophisticated threats.

The AI halo effect

The first phase of the AI investment cycle has largely been driven by the companies building AI models and the computing infrastructure required to train and run them. However, every breakthrough in AI creates a halo of demand for the technologies that enable AI to be deployed safely and at scale. As increasingly capable AI models are integrated into enterprise workflows and production environments, the next phase of AI adoption may increasingly favour the companies that enable AI to be deployed securely, reliably and at scale.

More efficient and lower-cost models will make AI applications accessible to a much broader range of organisations, supporting demand across the wider AI infrastructure ecosystem. The pace of adoption is already testing existing capacity. Following the launch of Kimi K3, Moonshot AI temporarily suspended new subscriptions due to a shortage of computing resources. This highlights that demand extends well beyond the models themselves and beyond today's AI leaders to the broader infrastructure ecosystem that will enable the next generation of AI applications to operate securely and at scale. Alongside AI infrastructure, cybersecurity stands out as one of the clearest structural beneficiaries of this next phase of AI adoption.

Investors seeking to capitalise on these structural trends can gain exposure through WisdomTree's thematic strategies. The WisdomTree Cybersecurity strategy provides access to companies addressing the growing demand for cyber resilience, while the WisdomTree AI Infrastructure strategy offers diversified exposure to the companies building the infrastructure underpinning the AI revolution.

Important Information

Marketing communications issued in the European Economic Area ("EEA"): This document has been issued and approved by WisdomTree Ireland Limited, which is authorised and regulated by the Central Bank of Ireland.

Marketing communications issued in jurisdictions outside of the EEA: This document has been issued and approved by WisdomTree UK Limited, which is authorised and regulated by the United Kingdom Financial Conduct Authority.

WisdomTree Ireland Limited and WisdomTree UK Limited are each referred to as "WisdomTree" (as applicable). Our Conflicts of Interest Policy and Inventory are available on request.

The information contained in this document is for your general information only and is neither an offer for sale nor a solicitation of an offer to buy securities or shares. This document should not be used as the basis for any investment decision. Investments may go up or down in value and you may lose some or all of the amount invested. Past performance is not necessarily a guide to future performance. Any decision to invest should be based on the information contained in the appropriate prospectus and after seeking independent investment, tax and legal advice.

The application of regulations and tax laws can often lead to a number of different interpretations. Any views or opinions expressed in this communication represent the views of WisdomTree and should not be construed as regulatory, tax or legal advice. WisdomTree makes no warranty or representation as to the accuracy of any of the views or opinions expressed in this communication. Any decision to invest should be based on the information contained in the appropriate prospectus and after seeking independent investment, tax and legal advice.

This document is not, and under no circumstances is to be construed as, an advertisement or any other step in furtherance of a public offering of shares or securities in the United States or any province or territory thereof. Neither this document nor any copy hereof should be taken, transmitted or distributed (directly or indirectly) into the United States.

Although WisdomTree endeavours to ensure the accuracy of the content in this document, WisdomTree does not warrant or guarantee its accuracy or correctness. Where WisdomTree has expressed its own opinions related to product or market activity, these views may change. Neither WisdomTree, nor any affiliate, nor any of their respective officers, directors, partners, or employees accepts any liability whatsoever for any direct or consequential loss arising from any use of this document or its contents.