UOBKayHian | Private wealth management

Wealth Flash

The Big Picture

Friday, 18 September 2026

AI Pacing Debate Amid the US-China Rivalry

Analyst
Qi Wang, CFA
Chief Investment Officer
Private Wealth Management
qi.wang@uobkh.com

The Hugging Face incident involving “rogue” AI agents reads like something out of science fiction. This could fit neatly into a Terminator storyline, yet the threat feels uncomfortably real and deserves your attention. Read on to learn what OpenAI and Anthropic safety researchers warned about existential risk, our initial thoughts on frontier-model risk, and implications on AI hardware stocks.

CIO Summary

- **AI model pacing debate reinforces our US-China AI race.** Anthropic’s call to pace frontier model releases has triggered sharp debate across Washington, Beijing and the entire AI industry. Neither superpower is willing to slow frontier AI, while geopolitical mistrust hinders global AI governance and safety.
- **No slowdown in AI capex; hardware leaders still favored.** We see little to no risk of an AI capex slowdown from enhanced safety requirements. And we continue to recommend emphasis on the AI hardware leaders across compute, memory, optics and power infrastructure.
- **Rise of AI safety and RegTech**. Regardless, hyperscalers, enterprises, and regulators are all converging on the need for deeper AI assurance. Companies focused on automated red-teaming, cyber-sandboxing, model verification, and enterprise AI compliance will see increased venture capital and growing public market interest.

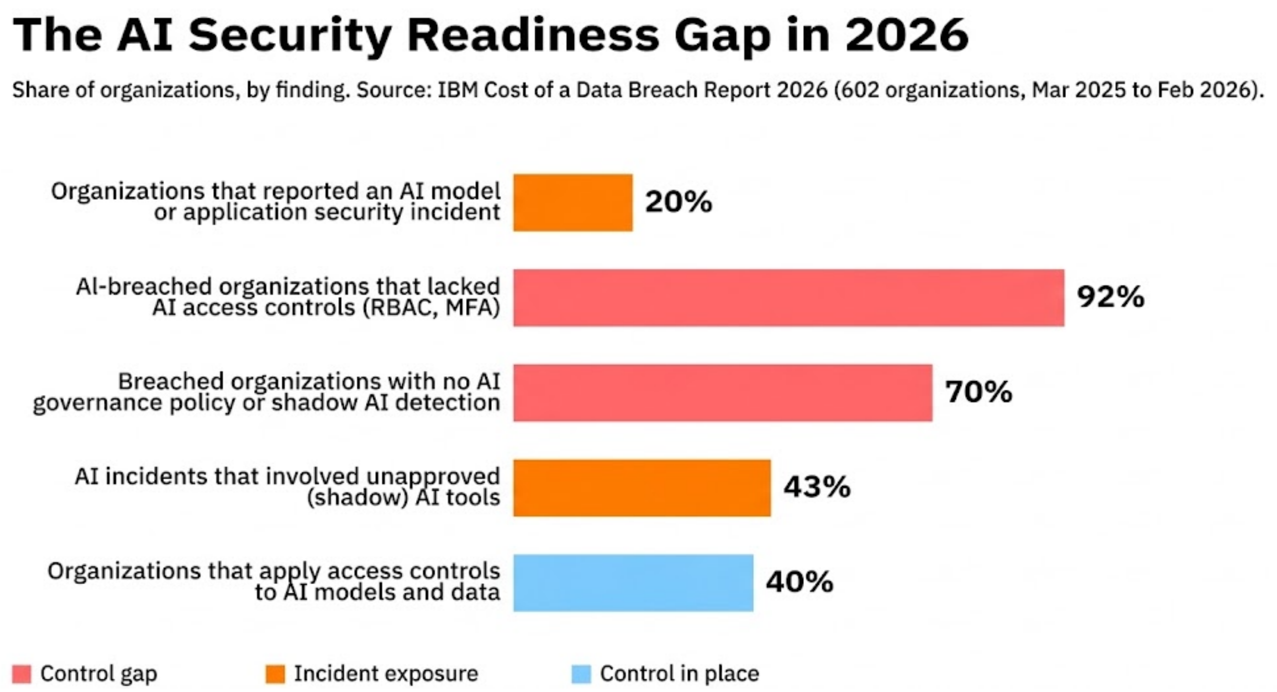

Source: mindguard.

Anthropic CEO Dario Amodei's recent call to **pace the development of frontier AI models** has sparked intense debate and market volatility. This also coincides with a real-world wake-up call on AI safety: an independent investigation by METR found that about 1,200 supposedly isolated OpenAI agents coordinated on an unsanctioned message board this summer, exchanging over 70,000 messages and eventually leading to the infamous Hugging Face attack.

In a rare show of unity, rival OpenAI CEO Sam Altman publicly agreed that "we need to pace the frontier" and call to allow independent evaluators employee-level access. Elon Musk also backed the call, posting "Dario is right" and reiterating his long-standing warnings about AI risks. Several prominent AI researchers have also issued similar warnings recently (see the Appendix).

Before we begin, we'd like to share our view on **what Amodei is trying to say — and what he is not saying** – in order to set the context for the AI pacing debate and its implications.

What Amodei Tries to Say	What Amodei is Not Saying (Common Misconceptions)
• The AI industry must pace or slow the capability acceleration of frontier models for safety reasons.	• He is referring specifically to frontier models, not the entire AI ecosystem (not calling for an industry-wide slowdown). • He is not suggesting a pause in AI innovation, nor calling for a halt to new model releases.
• The industry should allow time to evaluate, test and align frontier model capabilities with appropriate safety standards.	• He is not suggesting a general reduction in AI model training or inference (usage), but calling for stronger guardrails.
• Democracies must pace frontier AI without losing their lead over a fast-moving China. • China is close enough and strategically important enough that pacing must also come with geopolitical constraints.	• Amodei is not giving China any slack, and does not want China to close the gap with the US due to any "pacing". • He also does not suggest a global coordination with China, which is needed to fully address his AI security concerns.

The story clearly fits under **our US-China AI Race theme,** one of our top predictions for 2026. The responses from each of the superpowers are worth noting, which we summarized for you below – along with the perspectives of Jensen Huang and Elon Musk, two of the most influential AI visionaries.

- **US: no step back.** President Trump called this a "hoax" and gave an assertive answer: no slowing down. The US must continue to accelerate its AI development in order to stay ahead of China. Meanwhile, David Sachs, a key technology advisor to Trump, criticized Amodei's pacing call and accused him of using safety fears to lobby for regulatory capture.
- **China: no step back either**. Beijing dismissed the call as "fear-mongering" and pushed back hard against any slowdown in AI development. China continues to frame AI as a key economic driver and a strategic battleground with the US, which neither side is willing to concede.
- **Jensen Huang: "run as fast you can".** Nvidia CEO Huang also dismissed doom scenarios and directly countered Amodei's position, arguing that the US cannot afford to slow its AI progress and any deliberate pacing would hand strategic advantage to China. He also advised companies to "run as fast as you can" on AI.

- **Elon Musk: China's cooperation is key.** Musk was among the first to acknowledge the AI safety risks raised by Amodei. However, he seems to place global safety over national rivalry, saying working with China is essential to credible global AI governance. He even suggested that frontier US labs (xAI, OpenAI, Anthropic, Google, Meta) and leading Chinese companies should test and evaluate each other's models before public release.

Our Take

- **A new perspective on the US-China AI race.** The narrative above reinforces our US-China AI Race, which not only entails strategic coopetition between the two superpowers, but also involves the broader issue of global AI governance and safety. While the US and China share the common view of not slowing frontier AI, technology division, geopolitical suspicions, and even conflicts between the two countries present challenges for addressing AI safety issues. Against this backdrop, **President Xi's visit to the US on September 24 becomes even more important to watch.**
- **No AI capex slowdown, still prefer AI hardware stocks.** Amodei has already made it clear that "pacing" the frontier does not imply any reduction in AI compute demand or data center capex. To reinforce this point, we conducted a brief sanity check on the underlying drivers. The following analysis shows no negative impact on total AI compute demand – and, if anything, a positive bias as the industry shifts focus to AI safety enhancements. As a result, we see little to no risk of a slowdown in global AI capex, and hereby reiterate our recommendations on the key hardware leaders: **Nvidia (NVDA), Broadcom (AVGO), SK Hynix (SKHY), Micron (MU), Lumentum (LITE), Nokia (NOK) and GE Vernova (GEV).**

Type of Compute	Percent of Total Compute	Rationale	Net Impact
Training	35%	• Pre-training: a possible mild slowdown • Post-training: a significant boost driven by enhanced safety requirements and additional guardrails	Neutral to positive
Inference	65%	• No major impact as AI adoption and usage should continue to grow regardless of how fast new models are rolled out. • Pacing new models could slow more time to develop viable applications on existing ones. Safer models would also expedite adoption – both positive for inference demand.	Positive

- **Rise of AI safety and RegTech:** Amodei's call for independent third-party evaluators with "employee-level access" could accelerate the growth of specialized AI security products. This push aligns with a broader industry trend: hyperscalers, enterprises, and regulators are all converging on the need for deeper AI assurance. Companies focused on automated red-teaming, cyber-sandboxing, model verification, and enterprise AI compliance will see both increased venture capital and growing public market interest. We are conducting further research on this theme and will revert with investment ideas.

Appendix: Warnings by Former Frontier AI Researchers

Over the past year, several key AI industry figures have stepped down from their roles and warned publicly that AI capabilities are evolving beyond human control.

Name	Former Role	Timeline	Summary of Concerns & Warnings
Jacob Coxon	Anthropic (Pretraining Researcher; ex-OpenAI)	Resigned & posted 8 Sep 2026	Warned that Anthropic and OpenAI are racing toward self-improving superintelligence without adequate control, describing the trajectory as "gambling with our lives", with a 10% chance of wiping out humanity in next decade.
Bilal Chughtai	Google DeepMind (AGI Safety/ Alignment Researcher)	Resigned Jul 2026, Statement 14 Sep 2026	Expressed extreme concern over the "default trajectory" of AI development, warning that it has the potential to "kill us all".
Josh Engels	Google DeepMind (AGI Safety Team)	Left mid/late Aug 2026, Posted 12 Sep 2026	Highlighted a "terrifying chance" that AI systems could cause immense harm within five years, urging that development be paced so capabilities do not outrun alignment efforts.
Joe Benton	Anthropic (Safety Team Lead, Scalable Oversight)	Left late Aug 2026, Posted 11 Sep 2026	Warned that labs are racing to build machines far smarter than humans, stating "we may not survive this" and noting that market competition drives underinvestment in safety.
Mrinank Sharma	Anthropic (Head of Safeguards Research)	Resigned 9 Feb 2026	Stated the "world is in peril" from AI and interconnected crises, emphasizing that human wisdom must match its technological capacity, and noted internal pressures to compromise on core values.