5 May 2026

BY EMAIL ONLY

Attention: CII Owner & Board Chairman

Cybersecurity implications of frontier AI

I am writing to draw your attention to the cybersecurity implications of recent advances in frontier AI, and set out what we need to do in response.

In the past month, frontier AI has materially shifted the cybersecurity baseline for CIIs. On 7 April 2026, Anthropic announced Claude Mythos Preview, but restricted access to vetted defenders under Project Glasswing because of its advanced cyber capabilities. Anthropic stated that Mythos had already identified thousands of zero-day vulnerabilities. Shortly after, the UK AI Security Institute reported that Mythos could execute multi-stage attacks on vulnerable networks and had become the first model it tested to complete a 32-step end-to-end corporate network intrusion simulation, estimated to take an expert human around 20 hours. OpenAI's subsequent release of GPT-5.5 reinforces the same direction of travel: OpenAI assesses the model as having "High" cybersecurity capability under its Preparedness Framework, one step below "Critical".

These developments demand board-level and CEO attention, especially for CII owners and should not be left to IT departments. Frontier AI is accelerating at a rate where current assumptions in cyber risk management, on which your controls, measures and incident response plans were designed, may no longer be valid. Vulnerability discovery is becoming faster and cheaper. Social engineering is becoming more convincing and more personalised. Multi-stage attack chains can increasingly run without human intervention. Suppliers and interconnected systems face similarly heightened pressure. The window between vulnerability disclosure to system owners and exploitation by bad actors is narrowing, and the level of expertise required to mount a competent attack is falling.

To help CII Owners navigate the changed risk environment, CSA's alert (NCSC/Alert/2026/066) on 13 April 2026 has set out the immediate technical mitigations to be followed up on. Beyond this, we ask that Boards commission a review of whether your cybersecurity risk posture remains adequate in light of frontier AI development.

This review should consider:

1. whether the organisation's current cyber risk assessment takes relevant account of AI-enabled threats, both for IT and OT systems;

2. whether visibility over critical systems, internet-facing assets, privileged access, cloud services and third-party dependencies remains sufficient;

3. whether vulnerability management, patching, monitoring and incident response arrangements are fast enough as adversary tempo accelerates;

4. whether your organisation's own use of AI is appropriately governed, particularly where AI tools interact with sensitive data, software development, cybersecurity operations or critical systems; and

5. where AI can be used to augment current cybersecurity operations, including review of your organisation's code security.

This review should be tabled at the appropriate Board or executive governance risk committee. Where material gaps are identified, management should ensure that these are addressed through clear remediation plans and explicit risk acceptance decisions and where necessary, adjustments to cybersecurity investment priorities. CSA will engage your respective sector lead in the coming weeks to take stock of progress, understand challenges faced and discuss where we can work together to further strengthen your cybersecurity posture.

On our part, CSA will continue to monitor these developments, publish further technical guidance as the picture evolves, and work with our partners across industry and government to strengthen Singapore's collective cyber resilience.

Thank you for your continued partnership in safeguarding Singapore's essential services.

Yours sincerely,
David Koh
Commissioner of Cybersecurity