



C/2026/3795

16.7.2026

WARNING OF THE EUROPEAN SYSTEMIC RISK BOARD
of 25 June 2026
on systemic cyber risks stemming from frontier artificial intelligence models
(ESRB/2026/3)
(C/2026/3795)

THE GENERAL BOARD OF THE EUROPEAN SYSTEMIC RISK BOARD,

Having regard to the Treaty on the Functioning of the European Union,

Having regard to the Agreement on the European Economic Area ⁽¹⁾, and in particular Annex IX thereto,

Having regard to Regulation (EU) No 1092/2010 of the European Parliament and of the Council of 24 November 2010 on European Union macroprudential oversight of the financial system and establishing a European Systemic Risk Board ⁽²⁾, and in particular Article 3(2), point (c), and Articles 16 and 18 thereof,

Having regard to Decision ESRB/2011/1 of the European Systemic Risk Board of 20 January 2011 adopting the Rules of Procedure of the European Systemic Risk Board ⁽³⁾, and in particular Articles 18 and 19 thereof,

Whereas:

- (1) Leading artificial intelligence (AI) providers have developed Frontier AI Models (FAIMs), which are highly capable in the cybersecurity domain and able to carry out fully automated cyber-attacks on complex systems, including through vulnerability discovery and the development and weaponisation of exploits, thereby fundamentally changing the cyber-threat landscape.
- (2) FAIMs are significantly better than earlier AI models in finding ways to carry out cyberattacks, outperforming earlier models in cost, speed and accuracy, now rivalling leading human experts.
- (3) FAIMs have the capability to threaten the systems underpinning the information and communications technology (ICT) environments on which financial infrastructure relies, increasing the probability and severity of cyber incidents with systemic impact.
- (4) FAIMs have demonstrated the capability to discover vulnerabilities and craft novel exploits, thereby exposing major operating systems and software used across almost all organisations' ICT environments to the risk of cyber-attacks.
- (5) Historically, the discovery of such vulnerabilities has been rare and often made by security researchers, allowing for targeted remediation. In those cases, software providers are often given a standardised timeframe – usually 90 days – to patch the vulnerability before its public disclosure. However, FAIMs are likely to increase the volume of vulnerabilities found.
- (6) Current patching practices in the financial system are predominantly reactive, relying on periodic updates and ad hoc responses to disclosed vulnerabilities. This approach works in a threat environment where time allows for vulnerability discovery to remain manageable. However, with an increase in volume of vulnerabilities, current practices may become insufficient to maintain operational resilience.
- (7) This increase in volume also exposes current patching processes to overloading, as they prioritise remediation measures based on risk and require testing of software patches before implementation to ensure operational stability. If too many vulnerabilities with critical impact are discovered within a short timeframe so that the number of critical software patches required is substantially increased, financial institutions may have to decide between leaving themselves exposed to significant cyber risks, or reducing the patch-testing requirements and thereby risking operational incidents and outages.

⁽¹⁾ OJ L 1, 3.1.1994, p. 3, ELI: http://data.europa.eu/eli/agree_internation/1994/1/oj.

⁽²⁾ OJ L 331, 15.12.2010, p. 1, ELI: <http://data.europa.eu/eli/reg/2010/1092/oj>

⁽³⁾ OJ C 58, 24.2.2011, p. 4.

- (8) The crafting of weaponised exploits was previously largely done manually and took human experts days or weeks, whereas it can now be done by FAIMs in a matter of minutes or hours. This constitutes a collapse of defensive time buffers, which are needed to maintain the continuity of critical and important functions during remediation.
- (9) With the proliferation of weaponised exploits corresponding to critical vulnerabilities powered by AI and the corresponding reduction of defensive time buffers, the risk to individual critical or important functions and institutions increases sharply. If such risks materialise simultaneously across those functions and institutions, this could result in a permanent increase in systemic cyber risk for which, at present, there is no fully effective mitigation framework available. Due to the strong interconnectedness within the Union financial sector, as well as between that sector and other sectors and jurisdictions, this situation poses a systematic risk and has the potential to create systemic fragilities.
- (10) Furthermore, FAIMs are currently being developed by only a small number of AI providers, several of whom have expressed concerns that forthcoming FAIMs may be too powerful for unrestricted public access.
- (11) AI is already being used by malicious actors to enhance cyber-attacks. It is likely that such actors will eventually obtain, whether through leaks, theft, independent development, or open access, AI models with capabilities comparable to those currently deployed in controlled defensive settings.
- (12) While FAIMs may significantly increase offensive capabilities and the likelihood of sophisticated cyberattacks, they will also strengthen defensive capabilities, specifically regarding vulnerability discovery, data correlation and remedial action by individual institutions. However, in the short-to-medium term, the increase in speed, scale and accuracy of offensive capabilities are likely to outweigh the benefits.
- (13) The rapid development of FAIMs makes it necessary to address the identified risks to financial stability without delay so as to prevent their materialisation or at least to mitigate their potential impact. To ensure financial stability in the Union financial markets, all European Systemic Risk Board (ESRB) members should consider the implications of FAIMs on operational capability and resilience of financial institutions in the context of their macroprudential and microprudential policy actions. This is without prejudice to the monetary policy mandates of the central banks in the Union. As an example, and among other authorities' current initiatives, the ECB, in its role of banking supervisor, has requested the significant institutions to assess the impact of the evolving threat landscape without delay, and to develop by 31 October 2026 a comprehensive action plan outlining concrete measures aimed at addressing such risks ⁽⁴⁾.
- (14) As AI technology proliferates, it is important to ensure that all critical infrastructure and especially financial institutions are prepared to face cyber-attacks powered by FAIMs, and make timely and appropriate preparations. However, such defensive efforts by single entities would be insufficient. An effective response and risk mitigation requires a coordinated answer involving all affected parties including AI providers, software providers, security firms, open-source maintainers, financial institutions, and authorities at both national and Union level.
- (15) When addressing the issues above, it is important to recall relevant Union legal and policy frameworks, in particular Regulation (EU) 2022/2554 of the European Parliament and of the Council ⁽⁵⁾, Regulation (EU) 2024/1689 of the European Parliament and of the Council ⁽⁶⁾, Regulation (EU) 2024/2847 of the European Parliament and of the Council ⁽⁷⁾ and Recommendation of the European Systemic Risk Board (ESRB/2021/17) ⁽⁸⁾.

⁽⁴⁾ The letter addressed to significant institutions is available on the ECB's Banking Supervision website at www.bankingsupervision.europa.eu.

⁽⁵⁾ Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011 (OJ L 333, 27.12.2022, p. 1, ELI: <http://data.europa.eu/eli/reg/2022/2554/oj>).

⁽⁶⁾ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (OJ L, 2024/1689, 12.7.2024, ELI: <http://data.europa.eu/eli/reg/2024/1689/oj>).

⁽⁷⁾ Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements and amending Regulations (EU) No 168/2013 and (EU) 2019/1020 and Directive (EU) 2020/1828 (Cyber Resilience Act) (OJ L, 2024/2847, 20.11.2024, ELI: <http://data.europa.eu/eli/reg/2024/2847/oj>).

⁽⁸⁾ Recommendation of the European Systemic Risk Board of 2 December 2021 on a pan-European systemic cyber incident coordination framework for relevant authorities (ESRB/2021/17) (OJ C 134, 25.3.2022, p. 1).

- (16) Regulation (EU) 2022/2554 provides a comprehensive framework for financial institutions to identify, manage, monitor and mitigate ICT risks, including risks stemming from cyber-attacks. Given the interconnectedness of the financial sector and its reliance on ICT systems and third-party service providers, a successful cyber-attack affecting one or more financial institutions or to their critical service providers could have significant systemic implications. Financial supervisory authorities should give appropriate priority to supervisory activities in this area, having regard to the potential impact of such cyber-attacks on the stability, integrity and trustworthiness of the financial system, as well as to the operational, financial and consumer harm that may result.
- (17) Regulation (EU) 2024/1689 establishes harmonised rules for the placing on the market, putting into service and use of AI systems and general-purpose AI models in the Union, including, in certain circumstances, systems and models provided by AI providers established in third countries where the output produced by the AI system is used in the Union, or an AI model is integrated into a system placed on the Union market. This Regulation also provides specific rules, including a stricter regulatory regime, for models classified as general-purpose AI models with systemic risk.
- (18) Regulation (EU) 2024/2847 introduces mandatory cybersecurity requirements for manufacturers which cover the planning, design, development and maintenance of hardware and software products with digital elements placed on the Union market. It also requires manufacturers to handle vulnerabilities during the lifecycle of their products. Once fully applicable in December 2027, financial entities placing such products on the Union market will also need to ensure those products comply with these requirements.
- (19) Recommendation ESRB/2021/17 recommended the development, by the European Supervisory Authorities (ESAs) jointly through the Joint Committee and together with the European Central Bank (ECB), the ESRB and the relevant national authorities, of an effective Union-level coordinated response in the event of a major cross-border cyber incident or related threat that could have a systemic impact on the Union's financial sector. This resulted in the establishment of the pan-European Systemic Cyber Incident Coordination Framework (EU-SCICF), which provides a valuable and operational platform for information exchange and coordination among financial authorities.
- (20) Against this background, the ESRB hereby adopts a warning of a general nature addressing the significant financial stability risks amplified by the development of FAIMs. The warning takes account of the analytical work in the ESRB note Addressing Frontier AI Models with cyber capabilities from a financial stability perspective published in June 2026 ⁽⁹⁾,

HAS ADOPTED THIS WARNING:

SECTION 1

Warning

New risk landscape

The rapid development of Frontier AI Models (FAIMs) with cyber capability brings material changes to the risk landscape for the Union financial system. Current evidence indicates that FAIMs are capable of discovering vulnerabilities, generating working exploits and autonomously executing full-scale cyber-attacks at a speed, scale and level of accuracy far exceeding previous AI models. This constitutes a paradigm shift in the cybersecurity domain and an inflection point in terms of AI capability.

The ESRB considers these developments to be a source of systemic risks to the financial system, with the potential to generate significant adverse effects on cybersecurity through accelerated vulnerability discovery, reduction in time available for response and effective defence, increased scale and sophistication of cyber-attacks, and the introduction of new dependencies and concentration risks within the financial system. While the ESRB recognises the strengthened defensive capabilities that FAIMs will offer financial institutions, these capabilities are likely to gradually materialise over the next few years, requiring financial institutions to operate through a transitional period characterised by elevated risk, high uncertainty and a rapidly evolving technological landscape. Given the highly digitalised and interconnected nature of the financial system, such adverse effects on cybersecurity could rapidly propagate, affecting critical and important functions across financial institutions, ultimately leading to systemic disruption and loss of confidence in the financial system.

⁽⁹⁾ Available on the ESRB's website at www.esrb.europa.eu.

The emergent FAIMs increase inherent ICT risk, leading to weakened operational resilience across the financial sector, particularly across four areas: (a) time, including the ability to rapidly patch complex systems without causing operational failure, taking into account the collapse of defensive time buffers between vulnerability discovery and the weaponisation of exploits; (b) the capability of defenders, including their ability to conduct FAIM-assisted automated and manual security testing, as well as their ability to protect, detect and respond to threats stemming from adversarial FAIM use; (c) concentration, including dependencies on a limited number of AI providers, dependent in turn on cloud providers, open source components and other widely used software; and (d) the capability of authorities, including their calibration of expectations, stress testing and preparedness measures in accordance with the development of FAIMs, so that operational failures do not become a source of broader instability.

Moreover, establishing and maintaining visibility of AI deployment within any financial institution – including fully understanding where and how AI is deployed – is challenging, taking into account that it is used in customer interfaces, internal agents, business processes, and integration with third parties.

Systemic assessment

The developments identified may materially alter both the scale and structure of cyber risk, increasing both direct and indirect risks to financial stability. The FAIM capabilities may place existing resilience frameworks under substantial pressure, especially if incidents are widespread and occur simultaneously. Incidents may also propagate rapidly across financial institutions and markets if critical third-party providers, shared technological ecosystems, or widely used open source components are affected, increasing the risk of related disruptions with systemic implications.

FAIMs alter three sources of asymmetry in the financial sector: the first between jurisdictions, the second between defenders and attackers, and the third between better resourced and less well-equipped financial institutions.

First, the current geographical concentration of leading AI providers outside the Union leaves the Union exposed to strategic dependency and geopolitical risk. In order to mitigate the knowledge and technological asymmetry between jurisdictions, the adoption of measures to facilitate adequate and proportionate access for the Union and its Member States to newly developed FAIMs is needed. Such measures should be timely, coherent and efficient, both in terms of policy and implementation, and should adequately take account of different Union financial sectors, benefiting from the expertise of the corresponding ESAs. This will alleviate not only an asymmetry between the Union and third-country jurisdictions but also between financial institutions established in different Member States, preserving a level playing field within and beyond the Union and minimising systemic risk. Arrangements facilitating access to FAIMs only for certain categories of institutions or in selected jurisdictions would contribute to the fragmentation of the single market in finance, reducing its liquidity and depth.

Second, as regards the asymmetry between attackers and defenders, FAIMs reduce the price of admission for attackers. Threat actors are likely to increasingly rely on FAIMs to conduct the more technically advanced and labour-intensive parts of offensive operations, decreasing costs. Defenders, by contrast, are constrained by operational requirements, dependencies and regulatory obligations, limiting the effect and scale to which they can benefit from FAIMs in the short-to-medium term, ultimately requiring longer periods to adjust.

Third, there are differences in how financial institutions are resourced and equipped to cope with the challenges posed by FAIMs. These differences may be generated by costs that are fixed rather than scaled, by resource constraints, and by limited access to specialists. Where there is asymmetry in the ability to cope with such challenges, the weakest link of the chain may affect the stability of the system.

Implications

It is essential to ensure that systemically important payment and settlement systems and financial market infrastructures remain protected against vulnerabilities emerging from FAIM use and development. Public and private operators should thoroughly review and update their cybersecurity frameworks to take into account such vulnerabilities. Authorities in charge of supervision or oversight should adopt measures to ensure that systems are adequately protected. Cooperation between authorities and private financial institutions should be pursued whenever necessary to increase the security and resilience of the Union financial system.

In its 2020 report on systemic cyber risk⁽¹⁰⁾, the ESRB identified speed, scale and malign intention of cyber threats as possible sources of propagation of systemic risk. Authorities should be aware that whilst FAIMs represent a new source of propagation, they also amplify the existing previously identified sources, and should integrate this insight into their policy action. Considering all the financial stability risks stemming from operational exposure to cyber risk in the financial sector, it is important to ensure that concerns identified in this Warning are reflected in the supervisory and oversight work and policy stances adopted by the relevant authorities, within their respective mandates.

Relevant authorities, within their respective competences and specifically under the framework established by Regulation (EU) 2022/2554, should be aware of three previously mentioned sources of asymmetry introduced by FAIMs in the financial sector. To reduce these asymmetries, cooperation arrangements and sectoral coordination should be considered alongside existing testing frameworks⁽¹¹⁾. Financial authorities should also ensure that private financial institutions' boards are fully committed to mitigating FAIM-driven cyber risks, that clear governance and accountability frameworks are in place, and that timely responses are adequately planned, along with corresponding internal investments.

Several possible developments considered in this Warning – whether the systemic impact, the geopolitical risk, or the changed balance between attackers and defenders – have the potential to impact the functioning of the financial system and the trust therein. To continuously assess these, the ESRB will consider the need for reflecting such developments in risk assessments, testing frameworks, and supervisory expectations, as well as in future scenario development. The ESRB will also monitor the use and development of FAIMs and their impact on the financial sector from a systemic risk perspective. Moreover, the ESRB will reassess developments in each quarterly risk assessment at the General Board level and consider other actions when needed.

SECTION 2

Definitions

For the purposes of this Warning, the following definitions apply:

- (1) 'AI provider' means a provider as defined in Article 3, point (3), of Regulation (EU) 2024/1689;
- (2) 'Frontier AI Models (FAIMs)' means advanced general-purpose AI models capable of materially affecting offensive or defensive cyber operations;
- (3) 'cyber-attack' means a cyber-attack as defined in Article 1 of Council Regulation (EU) 2019/796⁽¹²⁾, including when originating from inside the Union;
- (4) 'vulnerability' means a vulnerability as defined in Article 3, point (16), of Regulation (EU) 2022/2554;
- (5) 'exploit' means a tool, technique or method used to take advantage of one or more vulnerabilities to gain unauthorised access to systems, to modify them, or to disrupt them, or otherwise to give rise to ICT risk or ICT-related incidents;
- (6) 'weaponisation' means the preparation or adaptation of exploits enabling the systematic or scalable exploitation of vulnerabilities;
- (7) 'critical or important function' means a critical or important function as defined in Article 3, point (22), of Regulation (EU) 2022/2554;
- (8) 'general-purpose AI model' means an AI model as defined in Article 3, point (63), of Regulation (EU) 2024/1689;
- (9) 'general-purpose AI model with systemic risk' means a general-purpose AI model classified as presenting systemic risk in accordance with Article 51(1) of Regulation (EU) 2024/1689;
- (10) 'ICT-related incident' means an incident as defined in Article 3, point (8), of Regulation (EU) 2022/2554;

⁽¹⁰⁾ See Systemic cyber risk, ESRB, February 2020, available on the ESRB's website at www.esrb.europa.eu.

⁽¹¹⁾ E.g. threat intelligence-based ethical red teaming (TIBER), threat-led penetration testing (TLPT), advanced red teaming (ART) and cyber resilience stress testing (CyRST).

⁽¹²⁾ Council Regulation (EU) 2019/796 of 17 May 2019 concerning restrictive measures against cyber-attacks threatening the Union or its Member States (OJ L 129 I, 17.5.2019, p. 1, ELI: <http://data.europa.eu/eli/reg/2019/796/oj>).

- (11) 'ICT risk' means a risk as defined in Article 3, point (5), of Regulation (EU) 2022/2554;
- (12) 'operational resilience' means digital operational resilience as defined in Article 3, point (1), of Regulation (EU) 2022/2554.

Done at Frankfurt am Main, 25 June 2026.

*The Head of the ESRB Secretariat,
on behalf of the General Board of the ESRB*
Francesco MAZZAFERRO
