

From Rules to Resilience

Assessing the EU’s 2026 Action Plan on
Cybersecurity and Artificial Intelligence

Andrea Paone

Strategic Intelligence & Emerging Risks
Rapid Strategic Assessment No. 01

Product category	Rapid Strategic Assessment
Version	1.1
Research cut-off	17 July 2026
Revision date	18 July 2026
ORCID	0009-0003-6194-948X
DOI	10.5281/zenodo.21427785
Licence	CC BY 4.0
Review status	Independent research paper; not peer reviewed

Contents

Publication Information	1
Abstract	1
Executive Summary	2
Key Judgments	3
1. Introduction: Europe's Rules-to-Capability Problem	4
2. Scope, Method and Analytical Framework	5
2.1 Product category and evaluative scope	5
2.2 Sources, evidence stages and source traceability	5
2.3 Additionality rules	5
2.4 Multidimensional readiness framework	6
2.5 Dynamic-source protocol	7
3. AI-Enabled Cyber Threats: Demonstrated Capability and Uncertainty	7
3.1 Social engineering and task acceleration	7
3.2 Multi-step operations and ICS limitations	7
3.3 New attack surfaces	8
3.4 Threat calibration	8
4. Europe's Pre-existing Operational Baseline	9
5. What the Action Plan Adds: Additionality and Readiness	11
5.1 European third-party model-evaluation capacity	11
5.2 Secure access to advanced systems	11
5.3 ENISA-JRC secure testing platform	11
5.4 Defensive AI and open-source security	12
5.5 Grand Challenge and industrial pipeline	12
5.6 AI Factories and Gigafactories	12
5.7 Consolidated readiness profile	13
6. Critical Infrastructure: Uneven Absorption Capacity	14
6.1 Finance	14
6.2 Electricity	15
6.3 Healthcare	15
6.4 Comparative absorption profile	15
6.5 Cross-sector transmission	15
7. Strategic Dependencies and Operational Constraints	16
7.1 Functional dependency profile	16
7.2 Models and cloud	16
7.3 Compute, software, suppliers and skills	17
8. Alternative Interpretations and Falsifiers	17
8.1 Operational convergence	17
8.2 Regulatory layering	18

8.3 Cross-cutting condition: capability growth with continued dependence 18

8.4 Orchestration and sequenced AI-specific operationalisation 18

9. Indicator Register for 2026–2028 19

10. Conclusion 20

Limitations 21

Appendix A. Measurement and Scoring Rules 21

 A.1 Readiness profile rules 21

 A.2 Missing and conflicting evidence 22

 A.3 Sensitivity check 22

Appendix B. Public Evidence Matrix 22

Appendix C. Source Version Note 23

About the Author 24

Bibliography 24

Publication Information

Author: Andrea Paone

Series: *Strategic Intelligence & Emerging Risks* — Publication No. 01

Product category: Rapid Strategic Assessment

Version: 1.1, revised 18 July 2026

Original publication: Version 1.0, 17 July 2026

Research cut-off: 17 July 2026

Correspondence: research@aletheia-technologies.it

Personal contact: andrea.pao@outlook.it

ORCID: <https://orcid.org/0009-0003-6194-948X>

Research portfolio: <https://aletheia-technologies.it/en/research/>

DOI: [10.5281/zenodo.21427785](https://doi.org/10.5281/zenodo.21427785)

Licence: Creative Commons Attribution 4.0 International (CC BY 4.0)

Editorial imprint: Aletheia Technologies — Strategic Intelligence & Research (*provisional name*)

Status: Independent Research Initiative; not a legally incorporated entity

Review status: Independent research paper; not peer reviewed

Recommended citation: Andrea Paone, *From Rules to Resilience: Assessing the EU's 2026 Action Plan on Cybersecurity and Artificial Intelligence*, *Strategic Intelligence & Emerging Risks*, Rapid Strategic Assessment No. 01, version 1.1 (18 July 2026), <https://doi.org/10.5281/zenodo.21427785>.

Revision note. Version 1.1 reclassifies the publication as a Rapid Strategic Assessment and strengthens the method, evidence traceability, source-version control, sector comparison, dependency analysis and monitoring framework. The central proposition and substantive research cut-off are unchanged.

The views and judgments expressed in this paper are those of the author. The editorial imprint is provisional and does not represent a legally incorporated company. Copyright © 2026 Andrea Paone. This work is licensed under the [Creative Commons Attribution 4.0 International \(CC BY 4.0\)](#).

Abstract

This Rapid Strategic Assessment provides an ex-ante assessment of the additionality, implementation readiness and capability-development pathways of the European Union's 2026 Action Plan on Cybersecurity and Artificial Intelligence. It examines whether the Plan can translate Europe's regulatory and institutional framework into operational capabilities against AI-enabled cyber threats, particularly threats affecting critical infrastructure. A source-traceable framework separates inherited mechanisms, measures redirected or reprioritised by the Plan, and genuinely Plan-specific initiatives. Readiness is assessed across mandate and governance, resources and implementation machinery, technical readiness, access and integration, and operational use or effects. Given the ten-day interval between publication and the research cut-off, the paper does not evaluate impact or attribute changes in European cyber performance to the Plan. It finds that the Union

entered July 2026 with substantive horizontal cyber capacity, while the Plan's distinctive AI-specific components remained largely prospective. The Plan's immediate value therefore lay in orchestration and implementation design. Sectoral absorption is likely to be uneven, and strategic dependencies are better assessed through access, auditability, continuity, operational control and substitutability than through ownership alone.

Contribution: descriptive, comparative and analytical.

Keywords: European Union; artificial intelligence; cybersecurity; critical infrastructure; ENISA; operational resilience; strategic autonomy.

Executive Summary

The European Union's 2026 Action Plan on Cybersecurity and Artificial Intelligence does not create a new EU-wide cybersecurity and resilience architecture from the ground up. It builds on pre-existing legal, institutional and operational mechanisms—including NIS2, the CSIRTs Network, EU-CyCLONe, the Cyber Solidarity Act, the European Cybersecurity Reserve, DORA, the Cyber Resilience Act, the AI Act and EuroHPC—and seeks to orient them towards risks and opportunities created by advanced AI.¹

At the research cut-off, only ten days after publication, no operational outcome could reasonably be attributed to the Plan. Its immediate contribution was primarily **orchestration and implementation design**: identifying gaps in independent model evaluation, access to advanced systems, secure testing for critical sectors, defensive vulnerability management and European AI capacity. The most distinctive AI-specific measures remained at announcement, mandate, planning or procurement stage.²

The assessment distinguishes four evidence stages: legal or policy design, institutionally reported implementation, independently verified implementation, and operational performance or outcomes. Official Commission and ENISA material is authoritative for mandates, dates and institutional status, but does not by itself establish service-level availability or effectiveness. This distinction is especially important where Commission pages use non-identical descriptions of AI Factory operability and where dynamic sources revise numeric results.³

¹ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>.

² European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>; European Commission, "Governance and Enforcement of the AI Act," last updated 7 July 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-act-governance-and-enforcement>.

³ European Commission, "AI Factories," last updated 23 April 2026, accessed 18 July 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>; European Commission, "Artificial Intelligence," last updated 3 June 2026, accessed 18 July 2026, https://commission.europa.eu/topics/artificial-intelligence_en; European Commission, "Commission Marks One Year of the AI Continent Action Plan with Two New Reports on AI Adoption and Policymaking," 9 April 2026, accessed 18 July 2026, <https://digital-strategy.ec.europa.eu/en/news/commission-marks-one-year-ai-continent-action-plan-two-new-reports-ai-adoption-and-policymaking>; UK AI Security Institute, "How Fast Is Autonomous AI Cyber Capability Advancing?," 13 May 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>; UK AI Security Institute, "Our Evaluation of OpenAI's GPT-5.5 Cyber Capabilities," 30 April 2026, subsequently updated, accessed 18 July 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-openais-gpt-5-5-cyber-capabilities>.

Finance, electricity and healthcare were selected to maximise variation across regulatory maturity, technical architecture and organisational fragmentation. Finance appears comparatively well positioned to absorb common capabilities; electricity combines strong governance with specialised OT/ICS constraints; healthcare presents high need but uneven local capacity. These are comparative analytical judgments, not official sector scores.

Overall assessment: the Plan establishes a credible pathway from regulation towards capability, but its specifically AI-cyber operational layer remained largely prospective as of 17 July 2026. The strongest public evidence concerned inherited horizontal mechanisms rather than new Plan-specific services. **Confidence: moderate.**

Key Judgments

1. **The EU's inherited horizontal cybersecurity mechanisms were more operationally developed than the Action Plan's new AI-specific measures.** Existing cooperation, crisis-management and sectoral-resilience frameworks provide a genuine base, but they were not created by the Plan. **Confidence: high.**
2. **The Plan's most credible short-term contribution is orchestration and sequencing rather than immediate capability creation.** It aligns institutions and creates implementation pathways for model evaluation, secure access, protected testing, vulnerability management and compute. **Confidence: moderate.**
3. **No observable operational outcome could yet be attributed to the Plan at the 17 July 2026 cut-off.** The ten-day interval makes causal evaluation impossible. **Confidence: high.**
4. **Critical sectors differ materially in their capacity to absorb common European capabilities.** Finance, electricity and healthcare vary across governance, human capability, technical integration, procurement and remediation capacity. **Confidence: high that current sector ecosystems differ materially; moderate regarding their future absorption of Plan-specific capabilities.**
5. **Public evidence demonstrates constrained multi-step autonomy in controlled cyber ranges, including completion of a range containing an OT/ICS-oriented attack chain, but not reliable attacks against defended, heterogeneous real-world industrial systems.** Conflicting and revised AISI run counts reinforce the need to treat benchmark figures as versioned measurements rather than fixed capability constants. **Confidence: moderate.**
6. **Technological dependencies constrain specific dimensions of European freedom of action without implying an absence of indigenous capacity.** Access, auditability, continuity, operational control and substitutability are more informative than ownership alone. **Confidence: moderate.**

7. **Every measure must be assessed at the evidence stage appropriate to its object type.** A legal framework, coordination network, response service, compute infrastructure and future platform should not be treated as directly commensurable merely because each can be assigned a single maturity label. **Confidence: high.**
-

1. Introduction: Europe's Rules-to-Capability Problem

By July 2026, the European Union had developed an extensive body of legislation and institutions covering network security, product security, crisis coordination, critical-entity resilience, financial-sector resilience and the governance of advanced AI. The relevant question was therefore no longer whether the Union possessed rules, but whether those rules and institutions could support effective operational action.

The strategic problem was whether this framework could keep pace with an evolving threat environment. Generative and agentic AI can already assist vulnerability research, social engineering, code production, reconnaissance and the automation of bounded technical tasks. The same technologies can support defenders through detection, triage, malware analysis and vulnerability prioritisation.⁴ Their effect will depend on who can access, evaluate, integrate and govern them most effectively.

The Commission published the *EU Action Plan on Cybersecurity and Artificial Intelligence* on 7 July 2026. It set out three objectives: promoting the safe and responsible use of advanced AI, reinforcing European cybersecurity and resilience, and scaling European AI capabilities for cybersecurity.⁵ The Plan proposed additional model-evaluation capacity, a blueprint for structured access to advanced systems, a secure ENISA–JRC testing platform for critical sectors, stronger use of AI in vulnerability management, the EU Grand Challenge on AI-assisted vulnerability remediation and continued investment in AI Factories and future Gigafactories.⁶

This paper assesses whether those measures translate regulation into operational capability, with particular attention to critical infrastructure. It argues that the Plan is best understood as an **orchestration and sequencing framework with prospective AI-specific operationalisation**: it builds on real European capabilities, while its distinctive AI-cyber mechanisms remained prospective or in early mobilisation at the cut-off.

⁴ ENISA, *ENISA Threat Landscape 2025*, version 1.2 (9 January 2026), 6, 8, 14–15, <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>; Andreas Happe, Aaron Kaplan, and Jürgen Cito, “LLMs as Hackers: Autonomous Linux Privilege Escalation Attacks,” *Empirical Software Engineering* 31, article 70 (2026), <https://doi.org/10.1007/s10664-025-10758-3>; UK AI Security Institute, “How Fast Is Autonomous AI Cyber Capability Advancing?,” 13 May 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>.

⁵ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>.

⁶ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>.

2. Scope, Method and Analytical Framework

2.1 Product category and evaluative scope

This publication is a **Rapid Strategic Assessment**. It provides an ex-ante assessment of policy additionality, implementation readiness and capability-development pathways. It is not an impact assessment and does not claim to measure changes in European cyber performance caused by the Plan.

The analysis covers the European Union, its institutions and agencies, and national or sectoral implementation where comparable public evidence is available. It focuses on the Commission, European AI Office, ENISA, Joint Research Centre, EuroHPC, CSIRTs Network, EU-CyCLONe and European Cybersecurity Reserve. The research cut-off is 17 July 2026.

Finance, electricity and healthcare are used as comparative critical-sector cases. They were selected to maximise variation across three dimensions: regulatory and supervisory maturity, technical architecture, and organisational fragmentation. The cases are illustrative rather than statistically representative; transport, water, digital infrastructure and ICT service management provide complementary evidence.

2.2 Sources, evidence stages and source traceability

The research uses structured document analysis, legislative and institutional mapping, a claim ledger, a public evidence matrix, comparative sector analysis and alternative-hypothesis testing. Primary legislation and official institutional material are prioritised for mandates, dates, budgets and formally reported implementation. Agency reports, peer-reviewed research and government technical benchmarks support assessment of sectoral conditions and the threat environment.

Evidence is classified by stage:

1. **Design evidence:** legislation, communications, mandates, plans and announced objectives;
2. **Institutional implementation evidence:** procurement, staffing, platform development, exercises or deployment reported by the responsible institution;
3. **Independent implementation evidence:** corroboration by auditors, users, operators, procurement records or independent technical assessments;
4. **Performance or outcome evidence:** repeated use, service availability, response times, remediation, operational effects or measurable improvement.

A source can be authoritative about design without proving performance. An announcement, funding allocation or institutionally reported status is therefore not treated as independent evidence of effectiveness. No public evidence was identified as of the cut-off for outcomes attributable to the Plan itself.

2.3 Additionality rules

To avoid attributing pre-existing capabilities to the Plan, measures are classified as:

- **Inherited:** operational or developing before 7 July 2026;

- **Redirected, repurposed or reprioritised:** a pre-existing mechanism is assigned an AI-cyber role or political priority without a demonstrated change in resources, scope or timing;
- **Accelerated:** public evidence shows an observable change in budget, staffing, procurement, timeline, scope or responsibility;
- **New and Plan-specific:** introduced directly by the Action Plan.

The term **accelerated** is not used solely because a pre-existing initiative is mentioned or politically reinforced by the Plan.

2.4 Multidimensional readiness framework

A single linear score can obscure differences between legal frameworks, networks, services, platforms and infrastructure. Version 1.1 therefore uses a five-dimension readiness profile for operational mechanisms and service-like measures:

Code	Dimension	0	1	2
G	Mandate and governance	no defined mandate	objective or responsible actor identified	mandate and governance formally established
R	Resources and implementation machinery	no public evidence	planning, tender or initial allocation	budget, procurement, staffing or delivery machinery in place
T	Technical readiness	concept only	prototype, pilot or partial capability	deployed technical capability or service
A	Access and integration	no defined access route	limited, test or partially specified access	intended users can access and integrate the capability
U	Operational use and effects	no exercise or use	exercise, isolated use or limited output	repeated use with observable outputs or outcomes

Profiles are written **G/R/T/A/U**. A descriptive summary band may be used for service-like measures only:

- **R0 — concept:** 0–2 total points;
- **R1 — mandated:** 3–4;
- **R2 — mobilising:** 5–6;
- **R3 — deployable:** 7–8;
- **R4 — operational evidence:** 9–10.

The summary band cannot override the dimension-level profile and is not applied to legal regimes or instruments whose object type makes aggregation misleading. Missing or conflicting evidence is not scored optimistically; ranges and confidence labels are used where necessary. A dimension-level score of 0 means that no qualifying public evidence was identified at the cut-off; it does not demonstrate that the capability or activity was absent.

2.5 Dynamic-source protocol

Material claims based on modifiable web pages are recorded with access date, page update date where available, the proposition observed and any later revision identified. Version 1.1 includes a Source Version Note and a supplementary Dynamic Source Register. Where two official pages conflict, the paper reports the divergence and narrows the permitted inference instead of selecting one description without explanation.

3. AI-Enabled Cyber Threats: Demonstrated Capability and Uncertainty

The Action Plan rests on a dual-use assessment. AI can strengthen defenders, but it can also lower the cost and increase the speed of selected offensive tasks. The evidence supports concern without supporting claims of general autonomous cyber warfare.

3.1 Social engineering and task acceleration

Generative models can produce persuasive and multilingual content, personalise messages and support iterative interaction. ENISA identifies AI as an optimisation tool for phishing and social engineering and as an emerging target through attacks on AI systems and supply chains.⁷ The principal change is one of scale, speed and accessibility. Attackers still require target data, delivery infrastructure, credential collection, technical access and a method of exploitation.

AI can also assist coding, vulnerability research, documentation analysis and reconnaissance. Peer-reviewed experiments on autonomous Linux privilege escalation show that model-based agents can succeed in bounded, controlled scenarios while exhibiting weaknesses in command validity, common-sense reasoning, error handling and multi-step exploitation, particularly where temporal dependencies are involved.⁸ Such results demonstrate meaningful component capability, not reliable operation against unknown and defended networks.

3.2 Multi-step operations and ICS limitations

The strongest public evidence of longer-horizon capability comes from the UK AI Security Institute's controlled cyber ranges. In an April 2026 evaluation, Claude Mythos Preview became the first model to complete *The Last Ones*—a 32-step simulated corporate-network attack—from beginning to end, succeeding in three of ten attempts and averaging 22 completed steps. AISI's 13 May update reported that a newer Mythos checkpoint completed *The Last Ones* in six of ten attempts and the previously unsolved *Cooling Tower* range in three of ten attempts. That update reported GPT-5.5 completing *The Last Ones* in three

⁷ European Union Agency for Cybersecurity (ENISA), *ENISA Threat Landscape 2025*, version 1.2 (9 January 2026), 6, 8, 14–15, <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>.

⁸ Andreas Happe, Aaron Kaplan, and Jürgen Cito, "LLMs as Hackers: Autonomous Linux Privilege Escalation Attacks," *Empirical Software Engineering* 31, article 70 (2026), <https://doi.org/10.1007/s10664-025-10758-3>.

of ten attempts; AISI's dedicated GPT-5.5 page reports two of ten after a grading correction. Version 1.1 records both official figures and does not treat the one-run difference as evidence of a different strategic trend.⁹

These results materially strengthen the evidence for constrained autonomous multi-step capability, but they do not establish reliable attacks against real-world critical infrastructure. AISI describes the ranges as small, vulnerable and undefended environments in which initial network access had already been obtained. They lacked active defenders and defensive tooling, and the agents incurred no penalty for actions that would trigger alerts. Completion rates remained inconsistent, and success on one controlled OT range in three of ten attempts does not demonstrate robust performance against heterogeneous, safety-constrained and actively defended industrial systems.¹⁰

OT environments contain proprietary protocols, long-lived equipment, safety constraints, segmented networks and operator-specific processes. AI may already help with reconnaissance, phishing of engineers, code adaptation or analysis of technical documentation. Converting access into a physical effect normally requires additional domain knowledge, persistence and interaction with the target environment.

3.3 New attack surfaces

AI systems themselves expand the attack surface. Agentic systems can combine planning, memory and tool use with access to logs, cloud controls or remediation systems. They may be exposed to prompt manipulation, memory or retrieval poisoning, malicious tool-output injection, unsafe tool use, model manipulation and supply-chain compromise.¹¹ A defensive agent may therefore become a privileged target.

Model evaluation is necessary but insufficient. Security must also cover the deployed system: integrations, permissions, data sources, monitoring, human oversight and fail-safe behaviour.

3.4 Threat calibration

The threat can be represented as a progression:

1. AI-assisted content and analysis;
2. automation of bounded tasks;

⁹ UK AI Security Institute, "How Fast Is Autonomous AI Cyber Capability Advancing?," 13 May 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>; UK AI Security Institute, "Our Evaluation of OpenAI's GPT-5.5 Cyber Capabilities," 30 April 2026, subsequently updated with a grading correction, accessed 18 July 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-openais-gpt-5-5-cyber-capabilities>; UK AI Security Institute, "Our Evaluation of Claude Mythos Preview's Cyber Capabilities," 13 April 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-claude-mythos-previews-cyber-capabilities>.

¹⁰ UK AI Security Institute, "Measuring AI Agents' Progress on Multi-Step Cyber Attack Scenarios," government technical benchmark, 16 March 2026, <https://www.aisi.gov.uk/research/measuring-ai-agents-progress-on-multi-step-cyber-attack-scenarios>; UK AI Security Institute, "Our Evaluation of Claude Mythos Preview's Cyber Capabilities," 13 April 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-claude-mythos-previews-cyber-capabilities>; UK AI Security Institute, "How Fast Is Autonomous AI Cyber Capability Advancing?," 13 May 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>.

¹¹ Nyashadzashe Tamuka et al., "Securing LLM-Based Agents against Cyberattacks: A Comprehensive Survey on Attack Techniques and Defense Strategies," *Journal of Computer Virology and Hacking Techniques* 22, article 38 (2026), <https://doi.org/10.1007/s11416-026-00622-3>; European Union Agency for Cybersecurity (ENISA), *ENISA Threat Landscape 2025*, version 1.2 (9 January 2026), 8, 14–15, <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>.

3. multi-step agentic assistance;
4. constrained autonomy in defined environments;
5. reliable end-to-end autonomy in unknown, defended environments.

Public evidence strongly supports the first two levels and increasingly supports the third. Evidence for the fourth now exists in defined, controlled ranges but remains inconsistent and context-specific. The fifth has not been demonstrated for complex, actively defended critical infrastructure in the reviewed public evidence.¹² Europe therefore has reason to build evaluation and preparedness while avoiding the presentation of controlled benchmark success as proof of reliable real-world autonomy.

4. Europe's Pre-existing Operational Baseline

The EU did not enter July 2026 with legislation but no operational architecture. It already possessed mechanisms for technical cooperation, crisis coordination, incident response, sectoral supervision, product-security reporting and AI computing.

The original NIS Directive established the CSIRTs Network. NIS2 subsequently strengthened that network and formalised EU-CyCLONe, which had been launched in 2020. ENISA supports both structures.¹³ Their existence provides a route for sharing and coordinating information on AI-enabled threats, although national capability and legal implementation remain uneven.

The Cyber Solidarity Act created the European Cybersecurity Alert System, a Cyber Emergency Mechanism and the European Cybersecurity Reserve.¹⁴ By July 2026, the Reserve had moved beyond a legislative commitment: ENISA administered the mechanism through a EUR 36 million contribution agreement, trusted providers had been selected, request procedures existed, the mechanism had been tested during Cyber Europe 2026, and it had first been deployed to Moldova in September 2025.¹⁵ ENISA continued to

¹² UK AI Security Institute, "How Fast Is Autonomous AI Cyber Capability Advancing?," 13 May 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>; UK AI Security Institute, "Our Evaluation of Claude Mythos Preview's Cyber Capabilities," 13 April 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-claude-mythos-previews-cyber-capabilities>; Andreas Happe, Aaron Kaplan, and Jürgen Cito, "LLMs as Hackers: Autonomous Linux Privilege Escalation Attacks," *Empirical Software Engineering* 31, article 70 (2026), <https://doi.org/10.1007/s10664-025-10758-3>.

¹³ Directive (EU) 2016/1148 of the European Parliament and of the Council of 6 July 2016 concerning measures for a high common level of security of network and information systems across the Union, art. 12, ELI: <https://data.europa.eu/eli/dir/2016/1148/oj>; Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, arts. 15–16, ELI: <https://data.europa.eu/eli/dir/2022/2555/oj>; ENISA, "EU CyCLONe," accessed 17 July 2026, <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cyclone>.

¹⁴ Regulation (EU) 2025/38 of the European Parliament and of the Council of 19 December 2024 laying down measures to strengthen solidarity and capacities in the Union to detect, prepare for and respond to cyber threats and incidents, ELI: <https://data.europa.eu/eli/reg/2025/38/oj>.

¹⁵ ENISA, "EU Cybersecurity Reserve," accessed 17 July 2026, <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cybersecurity-reserve>; ENISA, "The EU Cybersecurity Reserve - FAQ," accessed 17 July 2026, <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cybersecurity-reserve/the-eu-cybersecurity-reserve-faq>; ENISA, "Cyber Europe 2026: All Eyes on the EU's Collective Response and Resilience," 11 June 2026, <https://www.enisa.europa.eu/news/cyber-europe-2026-all-eyes-on-the-eus-collective-response-and-resilience>.

update the provider list and advertise new calls, so service coverage and specialised expertise were still developing.

DORA had applied since 17 January 2025 and provided finance with a structured regime for ICT risk, incident reporting, resilience testing and third-party oversight.¹⁶ It did not create AI-specific defence, but it supplied processes through which AI-enabled risk could be integrated.

The Critical Entities Resilience Directive provided the complementary physical-resilience layer, requiring national risk assessments, identification of critical entities and measures to maintain essential services.¹⁷ The Cyber Resilience Act created product-security and vulnerability-handling obligations, although its Single Reporting Platform was still being prepared for mandatory reporting from 11 September 2026.¹⁸

The AI Act required providers of general-purpose models presenting systemic risk to conduct evaluations and adversarial testing, mitigate systemic risks, report serious incidents and maintain cybersecurity protections.¹⁹ These obligations did not prove that authorities already possessed the personnel, compute and access necessary for independent verification—a gap recognised by the Action Plan.

Finally, EuroHPC and the AI Factories formed a growing general-purpose compute base. One Commission page described nineteen AI Factories and thirteen associated Antennas as operational and open to European users, while a separate Commission overview stated that work had begun on nineteen factories and that the majority were expected to be operational by the end of 2026.²⁰ The descriptions are not necessarily irreconcilable: network designation, partial service availability and site-level operationality can differ. This paper therefore separates network-level availability from service-level readiness and does not infer protected frontier-model evaluation, classified cyber testing or realistic OT/ICS capability from the general infrastructure label.

The baseline was therefore substantial but heterogeneous. It combined mature cooperation and crisis mechanisms with developing product-security tools and limited AI-specific technical capacity. That distinction is essential: the Action Plan inherited these capabilities; it did not create them.

¹⁶ Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector, ELI: <https://data.europa.eu/eli/reg/2022/2554/oj>.

¹⁷ Directive (EU) 2022/2557 of the European Parliament and of the Council of 14 December 2022 on the resilience of critical entities, ELI: <https://data.europa.eu/eli/dir/2022/2557/oj>.

¹⁸ 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, ELI: <https://data.europa.eu/eli/reg/2024/2847/oj>; ENISA, “Single Reporting Platform (SRP),” accessed 17 July 2026, <https://www.enisa.europa.eu/topics/product-security-and-certification/single-reporting-platform-srp>.

¹⁹ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, especially arts. 51–56, ELI: <https://data.europa.eu/eli/reg/2024/1689/oj>.

²⁰ European Commission, “AI Factories,” last updated 23 April 2026, accessed 18 July 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>; European Commission, “Artificial Intelligence,” last updated 3 June 2026, accessed 18 July 2026, https://commission.europa.eu/topics/artificial-intelligence_en.

5. What the Action Plan Adds: Additionality and Readiness

The Plan's additionality lies in connecting existing systems to new AI-specific functions. The evidence profile is asymmetric: inherited horizontal mechanisms show implementation or use, while Plan-specific measures are concentrated in design and early mobilisation.

5.1 European third-party model-evaluation capacity

The AI Act already established provider obligations, and the AI Office had launched a EUR 9 million tender in 2025 for technical support covering systemic-risk evaluation, including cyber-offence and agentic risks.²¹ The Plan's added value is a commitment to strengthen European third-party assessment capacity, expected by the Commission to become operational in 2027.²²

At the cut-off, the responsible institutional direction and procurement pathway were visible, but no deployed European service or completed evaluation output was publicly observable.

Additionality: new and Plan-specific expansion of an existing regulatory function.

Readiness profile: 2/1/0/0/0 (G/R/T/A/U), R1 — mandated.

5.2 Secure access to advanced systems

The proposed blueprint with ENISA recognises that authorities, researchers and operators may need access beyond ordinary commercial interfaces to evaluate advanced systems.²³ Relevant access could include documentation, logs, controlled environments, higher usage limits or provider cooperation. The Plan did not resolve eligible users, technical depth, liability, confidentiality or provider participation.

Additionality: new and Plan-specific.

Readiness profile: 1/0/0/0/0, R0 — concept.

5.3 ENISA-JRC secure testing platform

The planned platform would allow organisations in energy, transport, health, finance and public administration to test AI cybersecurity applications in secure simulated environments.²⁴ Its operational value will depend on realistic sectoral environments, protected data, OT/ICS components, service-level governance and clear onboarding arrangements. No accessible platform, service catalogue or completed test programme was identified.

²¹ European Commission, "EU AI Office Launches €9 Million Tender for Technical Support on GPAI Safety," 10 July 2025, <https://digital-strategy.ec.europa.eu/en/funding/eu-ai-office-launches-eu9-million-tender-technical-support-gpai-safety>.

²² European Commission, "Governance and Enforcement of the AI Act," last updated 7 July 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-act-governance-and-enforcement>.

²³ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>.

²⁴ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>.

Additionality: new and Plan-specific.

Readiness profile: 2/0/0/0/0, R0 — concept with assigned governance.

5.4 Defensive AI and open-source security

The Plan encourages AI-assisted vulnerability detection and remediation and assigns ENISA a supporting role through guidance, recommendations and work on critical open-source software.²⁵ These actions connect existing NIS2, CRA and vulnerability-management processes to AI-enabled practices. Public evidence reviewed for this assessment did not demonstrate an accelerated timeline, additional budget, expanded staffing or a new common defensive service.

Additionality: redirected or reprioritised, not demonstrably accelerated.

Readiness: varies with the underlying organisation and programme; no single aggregate band is assigned.

5.5 Grand Challenge and industrial pipeline

The proposed EU Grand Challenge on AI-assisted vulnerability remediation is intended to develop and scale European AI-powered cybersecurity solutions.²⁶ At the cut-off, the reviewed public record did not specify budget, rules, outputs, procurement route or adoption mechanism. An innovation competition may expand supply but is not itself an operational defensive capability.

Additionality: new and Plan-specific.

Readiness profile: 1/0/0/0/0, R0 — concept.

5.6 AI Factories and Gigafactories

AI Factories were inherited enabling infrastructure, not a creation of the cybersecurity Action Plan. Official Commission pages used non-identical descriptions of their operational status. General-purpose compute and support services may be available at network level without demonstrating that every site, service or cyber-specific function is operational.²⁷

For general AI Factory services, public evidence supports a profile of **2/2/2/1–2/0 (R3)**. Commission material supports network-level technical status and an access route, but the reviewed public record did not identify an exercise, actual user output or repeated operational use sufficient for a positive U score. For the protected cyber-specific functions contemplated by the 2026 Plan, the same infrastructure cannot be credited automatically; no separate operational profile is assigned without service-level evidence.

Future Gigafactories were intended to provide frontier-scale capacity. A dated Commission update referred to 76 responses to the expression-of-interest process, while the

²⁵ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>.

²⁶ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>.

²⁷ European Commission, “AI Factories,” last updated 23 April 2026, accessed 18 July 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>; European Commission, “Artificial Intelligence,” last updated 3 June 2026, accessed 18 July 2026, https://commission.europa.eu/topics/artificial-intelligence_en.

undated dedicated project page reported 77 proposals across 16 Member States and 60 sites.²⁸ The one-unit difference does not change the substantive assessment: the formal call and construction remained future steps.

Additionality: inherited enabling initiatives, politically reinforced and reprioritised; no acceleration claim is made without evidence of changed resources or timeline attributable to this Plan.

Gigafactory readiness profile: 2/1/0/0/0, R1 — mandated; mobilisation underway.

5.7 Consolidated readiness profile

The table applies the five-dimension framework only to operational mechanisms, services and service-like initiatives. Legal frameworks such as DORA, NIS2, the AI Act and the CRA remain essential baseline enablers but are not aggregated into the same score as platforms or response services.

Measure	Object type	Additionality	G/R/T/A/U	Summary	Confidence
CSIRTs Network	coordination network	inherited	2/2/2/2/1	R4	high on existence; moderate on outcomes
EU-CyCLONe	crisis network	inherited	2/2/2/1/1	R3	moderate
Cybersecurity Reserve	response service	inherited	2/2/2/2/1	R4	high on deployment; moderate on coverage
CRA Single Reporting Platform	reporting service	inherited, developing	2/2/1/0/0	R2	high
AI Factories, general services	compute/service network	inherited	2/2/2/1-2/0	R3	moderate due status and unit-of-analysis divergence
Third-party model-evaluation capacity	evaluation service	new	2/1/0/0/0	R1	moderate
Secure-access blueprint	governance instrument	new	1/0/0/0/0	R0	high
ENISA-JRC testing platform	testing service	new	2/0/0/0/0	R0	high
Grand Challenge	innovation instrument	new	1/0/0/0/0	R0	moderate
Gigafactories	infrastructure programme	inherited / reprioritised	2/1/0/0/0	R1	moderate

Interpretation: the profile does not imply that a coordination network and a compute service are substantively identical. It exposes which readiness dimensions are evidenced

²⁸ European Commission, “Commission Marks One Year of the AI Continent Action Plan with Two New Reports on AI Adoption and Policymaking,” 9 April 2026, accessed 18 July 2026, <https://digital-strategy.ec.europa.eu/en/news/commission-marks-one-year-ai-continent-action-plan-two-new-reports-ai-adoption-and-policymaking>; European Commission, “AI Gigafactories,” accessed 18 July 2026, https://commission.europa.eu/topics/competitiveness/competitiveness-coordination-tool-projects/ai-gigafactories_en.

and prevents a high legal or institutional status from being treated as proof of technical access, use or effects.

The central finding remains unchanged: Europe's strongest publicly evidenced assets were inherited horizontal mechanisms. The Plan's distinctive AI-specific components were concentrated in design and early mobilisation.

6. Critical Infrastructure: Uneven Absorption Capacity

ENISA's 2026 NIS360 assessment provides a comparative baseline, but it evaluates sectoral ecosystems rather than every operator.²⁹ The three cases were selected to maximise variation across regulatory maturity, technical architecture and organisational fragmentation:

- **finance:** high regulatory intensity and digital integration;
- **electricity:** mature governance combined with specialised OT/ICS and safety constraints;
- **healthcare:** high criticality, fragmented organisations and uneven resources.

Absorption capacity is assessed qualitatively across five dimensions: governance and supervision (**G**), human capability (**H**), technical integration (**T**), procurement and resources (**P**), and remediation or operational follow-through (**R**). Ratings are author assessments using **high, moderate or low** bands; they are not official ENISA scores.

The decision anchors are dimension-specific. **High** denotes consistent evidence of established sector-wide structures or capability, while recognising variation among organisations. **Moderate** denotes material capability combined with uneven coverage, persistent constraints or mixed evidence. **Low** denotes structural constraints likely to impede adoption or follow-through. Missing evidence is recorded as **not assessed**, not as low; conflicting evidence lowers confidence. Cell-level sources, uncertainty and confidence are documented in the supplementary Measurement and Scoring Register.

6.1 Finance

ENISA assessed banking as a high-maturity sector and reported that financial market infrastructures had moved into the high-maturity band, with DORA contributing to more structured risk management and preparedness.³⁰ Finance is therefore comparatively well positioned to integrate model evaluation, testing and AI-enabled threat information. The advantage is uneven by organisation size and constrained by concentration in cloud, identity, payment, software and shared providers.

²⁹ ENISA, *ENISA NIS360*, 5–13, <https://doi.org/10.2824/8928447>.

³⁰ ENISA, *ENISA NIS360*, 10, 12, 28–30, 55–57.

6.2 Electricity

Electricity was assessed as the most mature energy subsector, supported by established governance and sector-specific requirements.³¹ Its main constraint is technical: long-lived OT assets, strict availability and safety requirements, proprietary protocols and specialised engineering processes limit direct transfer from enterprise cyber tools. The ENISA–JRC platform would have high potential value only if it provides realistic industrial scenarios and combined engineering-cyber expertise.

6.3 Healthcare

Healthcare remained in the moderate maturity band and within the NIS360 risk zone despite progress.³² The sector combines high criticality with fragmented organisations, legacy technology, connected devices, third-party dependence and constrained resources. Shared European services may reduce the burden, but local personnel, procurement flexibility and protected testing environments remain necessary for adoption and remediation.

6.4 Comparative absorption profile

Sector	G/H/T/P/R profile	Overall assessment	Principal constraint	Confidence
Finance	H/H/H/H/H–M	comparatively strong, uneven by size	shared-provider concentration	moderate
Electricity	H/M/M/H/M	institutionally strong, technically specialised	OT/ICS safety, legacy and integration	moderate
Healthcare	M/L–M/L–M/L–M–M	highly uneven and resource-constrained	fragmentation, legacy and staffing	moderate

The profile is a structured comparative inference, not a measured index. Public evidence did not support organisation-level distributions or common quantitative thresholds across the three sectors.

6.5 Cross-sector transmission

Finance, electricity and healthcare depend on telecommunications, cloud, data centres, managed ICT and security providers, identity services, software supply chains and electricity.³³ A compromise in a common provider may propagate across sectors. EU-level capability must therefore pass through national authorities, sectoral supervisors, suppliers and operator systems; failure at any interface can leave a formally available European mechanism with limited effect on the ground.

³¹ ENISA, *ENISA NIS360*, 10, 12, 18–20, 46–48.

³² ENISA, *ENISA NIS360*, 5, 10, 12–13, 31–33, 58–60.

³³ ENISA, *ENISA NIS360*, 5–13, <https://doi.org/10.2824/8928447>; ENISA, *ENISA Threat Landscape 2025*, version 1.2 (9 January 2026), 6, 8, 14–15, <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>.

7. Strategic Dependencies and Operational Constraints

Technological dependence should be assessed by function rather than nationality alone. Version 1.1 applies five dimensions: access (**A**), auditability (**Au**), continuity (**C**), operational control (**O**) and substitutability (**S**). Ratings are qualitative author assessments based on the reviewed public record; no harmonised dataset exists.

For each dimension, **high** indicates comparatively strong and evidenced access, control, continuity or substitution capacity; **moderate** indicates conditional, mixed or contract-dependent capacity; **low** indicates a material constraint or an untested ability to maintain the function under stress. Missing evidence is recorded as **not assessed**, and conflicts lower confidence. The profile is a structured diagnostic rather than a measured index; cell-level confidence is generally low to moderate and is recorded in the supplementary Measurement and Scoring Register.

7.1 Functional dependency profile

Dependency	A/Au/C/O/S profile	Residual risk	Principal evidence gap
Frontier models	L/L/M/L/L	high	depth and continuity of evaluator access
Cloud infrastructure	M/M/M/M/M	medium-high	contract-level audit, exit and continuity testing
Advanced accelerators	L/L/M/L/L	high	supply continuity and practical substitution
Security software and open source	H/M/M/M–H/M	medium	component provenance and maintenance concentration
Managed security services	H/M/M/M/L–M	medium-high	provider concentration and cross-sector surge capacity

Legend: H = comparatively strong; M = conditional or mixed; L = constrained. A low substitutability rating does not imply that no alternative exists; it indicates that switching under operational stress may be slow, costly or technically disruptive.

7.2 Models and cloud

The access blueprint indicates that ordinary commercial interfaces may be insufficient for third-party model evaluation. Authorities may require protected environments, documentation, logs or provider cooperation. Diversification can reduce concentration, but models are not automatically interchangeable.

Cloud dependence combines jurisdiction, data control, privileged access, software supply chain and interruption risk. The Commission's 2026 procurement awarded four contracts worth up to EUR 180 million and used graduated sovereignty criteria. One consortium used Google Cloud-based technology through S3NS while services were to be op-

erated by EU companies.³⁴ The case supports a functional rather than binary approach to sovereignty, but it does not by itself demonstrate tested continuity or substitutability under stress.

7.3 Compute, software, suppliers and skills

AI Factories provide a developing base, but frontier evaluation also depends on advanced accelerators, networking, energy, software and access to target models. Gigafactories may reduce scale constraints later; they did not constitute available frontier-scale capacity at the cut-off. The Commission's technology-sovereignty package linked semiconductors, AI, cloud and open source, confirming their interdependence.³⁵

ENISA's 2025 investment survey identified recruitment and retention difficulties, patching delays and greater reliance on outsourced ICT and security services.³⁶ Shared European services can compensate for shortages, but organisations still need expertise to request assistance, interpret results and remediate findings. The Cybersecurity Reserve demonstrates that private capability can be operationalised through pre-arranged contracts and activation procedures; concentration nevertheless creates common-mode risk.

Europe therefore possesses substantial indigenous capacity while remaining constrained in frontier-model access, advanced processors, cloud technology, software components and specialist services. The relevant strategic objective is mission-relevant control and continuity, not complete autarky.

8. Alternative Interpretations and Falsifiers

Four interpretations were tested against the evidence. They are not assumed to be mutually exclusive: regulatory layering and operational convergence are competing descriptions of implementation, while capability growth with continued dependence is treated as a cross-cutting trajectory. Each is linked to an observable prediction and a condition that would weaken or falsify it.

8.1 Operational convergence

Prediction by 2028: joint exercises, interoperable services, shared procedures and measurable incident-response outputs connect the AI Office, ENISA, JRC, EuroHPC and existing crisis mechanisms.

Evidence at cut-off: complementary institutional roles were assigned in the Plan.

Evidence against: no Plan-attributable shared service, joint exercise or outcome was

³⁴ European Commission, "Commission Advances Cloud Sovereignty through Strategic Procurement," 17 April 2026, https://commission.europa.eu/news-and-media/news/commission-advances-cloud-sovereignty-through-strategic-procurement-2026-04-17_en; European Commission, "Sovereign Cloud Framework Explained," 1 June 2026, https://commission.europa.eu/news-and-media/news/sovereign-cloud-framework-explained-2026-06-01_en.

³⁵ European Commission, "Strengthening Europe's Tech Sovereignty," 3 June 2026, https://commission.europa.eu/news-and-media/news/strengthening-europes-tech-sovereignty-2026-06-03_en.

³⁶ ENISA, *NIS Investments 2025* (8 December 2025), 3–5, 15–16, 21, 30–31, <https://doi.org/10.2824/7442427>.

publicly observable.

Falsifier: persistent separation between agencies and no operationally used cross-institutional service.

Current assessment: plausible design hypothesis; not demonstrated operationally.

8.2 Regulatory layering

Prediction by 2028: new guidance, calls and obligations accumulate without material gains in access, service use, response time or remediation.

Evidence at cut-off: unresolved access and testing details, uneven national implementation and multiple future deliverables.

Evidence against: the EU already possessed usable crisis and response mechanisms, and the proposed functions address identifiable gaps.

Falsifier: repeated use of Plan-specific services with measurable operational outputs.

Current assessment: credible risk, but incomplete as a description of the existing baseline.

8.3 Cross-cutting condition: capability growth with continued dependence

Prediction by 2028: European evaluation, compute and response capacity expands while key dependencies in models, accelerators, cloud and specialist services remain.

Evidence at cut-off: expanding infrastructure and procurement coexisted with constrained access and substitutability.

Evidence against: credible exit options, diversified supply and demonstrated continuity would reduce the thesis.

Falsifier: operational substitution during a supplier disruption without material loss of mission capability.

Current assessment: strongly supported as a complementary trajectory rather than a mutually exclusive alternative; it is compatible with either successful orchestration or regulatory layering.

8.4 Orchestration and sequenced AI-specific operationalisation

Prediction by 2028: measures progress unevenly from mandate to procurement, limited access, pilot use and later operational outputs, with substantial sectoral differences.

Evidence at cut-off: inherited mechanisms were usable, while Plan-specific services remained prospective.

Evidence against: rapid deployment and repeated use of the new services would weaken the assessment that operationalisation remained prospective at the cut-off.

Falsifier: either no meaningful implementation pipeline emerges, or Plan-specific services become routinely accessible and used at European scale with observable outputs and credible links to improved remediation or response.

Current assessment: best fit to the evidence at the cut-off.

9. Indicator Register for 2026–2028

The thresholds below are analytical verification thresholds, not official EU targets. They specify what evidence would be sufficient to update the assessment. The nearest meaningful review point is **17 January 2027**, after the Q4 2026 deliverables should be observable; a second review is set for **17 July 2027** for measures expected to mature during 2027.

Indicator	Baseline at 17 July 2026	Verification threshold	Data owner / source	Frequency	Review date
European model-evaluation capacity	no deployed service or completed output identified	published method, responsible delivery structure and at least one completed third-party evaluation	AI Office / contractor	semi-annual	17 Jul 2027
Secure-access blueprint	not published	adopted blueprint defining eligible users, access depth, safeguards and contingency arrangements	Commission / ENISA	quarterly check	17 Jan 2027
ENISA–JRC testing platform	no accessible platform or catalogue	public service description, onboarding route and at least one enterprise and one OT/ICS-relevant environment	ENISA / JRC	quarterly	17 Jan 2027
Grand Challenge	concept without public rules or adoption route	published rules, budget, outputs and pathway from prototype to procurement or deployment	Commission	quarterly	17 Jan 2027
Cybersecurity Reserve	one reported deployment and one major exercise	recurring reporting on activations, sectors, response time and service coverage	ENISA	annual / per activation	17 Jul 2027
CRA reporting platform	reporting not yet active	operational reporting with published volume, triage or handling statistics	ENISA	quarterly after launch	17 Jan 2027
Sectoral absorption	no common participation and remediation dataset	sector-disaggregated evidence of onboarding, tests, exercises and remediation decisions	ENISA, national authorities, operators	annual	17 Jul 2027

Indicator	Baseline at 17 July 2026	Verification threshold	Data owner / source	Frequency	Review date
Protected cyber compute	no public allocation identified	documented compute allocation and access rules for protected model or cyber evaluation	EuroHPC / AI Office	semi-annual	17 Jul 2027
Functional dependency management	no harmonised public register	documented exit, continuity or substitution tests for material suppliers	Commission / operators	annual	17 Jul 2027

A future update should record not only whether an initiative was launched, but also who accessed it, what was tested, what remediation followed and whether service performance improved.

10. Conclusion

The EU’s 2026 Action Plan addresses a genuine strategic problem. Advanced AI is increasing the speed and accessibility of selected cyber tasks, while critical-sector defenders face persistent shortages, legacy systems and supplier dependencies.³⁷ Europe also possesses a substantial regulatory and operational baseline that can be adapted to the new environment.

The Plan had not produced new AI-cyber capabilities at European scale by the research cut-off. The Cybersecurity Reserve, CSIRTs Network, EU-CyCLONe, DORA, AI Office and AI Factories largely pre-dated it.³⁸ Its distinctive measures—expanded third-party model evaluation, structured access, secure testing and an AI-assisted vulnerability-remediation challenge—were still concentrated in design and early mobilisation.

The Plan’s immediate contribution was therefore to organise a capability-development pathway. The multidimensional readiness profile shows why no single maturity label is sufficient: mandates can be clear while resources, technical services, user access and operational effects remain absent. The same distinction applies to AI Factories, where network-level operationality does not establish the availability of every site, service or protected cyber function.

³⁷ UK AI Security Institute, “How Fast Is Autonomous AI Cyber Capability Advancing?,” 13 May 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>; ENISA, *NIS Investments 2025* (8 December 2025), 3–5, 15–16, 21, 30–31, <https://doi.org/10.2824/7442427>; ENISA, *ENISA NIS360*, 5–13, <https://doi.org/10.2824/8928447>.

³⁸ Directive (EU) 2022/2555 (NIS2), arts. 15–16, ELI: <https://data.europa.eu/eli/dir/2022/2555/oj>; Regulation (EU) 2025/38 (Cyber Solidarity Act), ELI: <https://data.europa.eu/eli/reg/2025/38/oj>; Regulation (EU) 2022/2554 (DORA), ELI: <https://data.europa.eu/eli/reg/2022/2554/oj>; Regulation (EU) 2024/1689 (Artificial Intelligence Act), especially arts. 51–56, ELI: <https://data.europa.eu/eli/reg/2024/1689/oj>; European Commission, “AI Factories,” last updated 23 April 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>.

Implementation risk is likely to arise from divergence. European institutions, national authorities, sectors and suppliers may progress at different speeds. Finance may integrate common capabilities through mature supervision, electricity requires specialised OT testing, and healthcare may struggle with local capacity. Strategic dependencies can create common points of failure even where European governance and infrastructure expand.

The central proposition remains that the Plan is an **orchestration and sequencing framework with prospective AI-specific operationalisation**. It is a credible architecture for capability development, not proof that Europe has already closed its AI-cyber gap. That judgment is falsifiable through the Indicator Register: it should be revised if the new services become accessible, repeatedly used and associated with observable remediation or performance outcomes.

Limitations

This assessment relies on public sources and cannot observe classified, internal or commercially confidential activities. Absence of public evidence is not evidence that an internal capability does not exist. National implementation and transparency vary. The ten-day interval between publication and cut-off prevents causal evaluation. NIS360 assesses sector ecosystems rather than individual operators. Experimental AI benchmarks may not transfer to live, defended infrastructure, and official benchmark pages may revise individual run counts. The readiness, absorption and dependency profiles are structured author assessments based on stated rules; they are not official EU ratings and do not substitute for organisation-level data. Dynamic pages could change after the cut-off; identified divergences are recorded in Appendix C and the supplementary Dynamic Source Register.

Appendix A. Measurement and Scoring Rules

A.1 Readiness profile rules

The five-dimension G/R/T/A/U framework applies only to operational mechanisms and service-like initiatives. Each dimension is scored 0–2 using the decision rules in section 2.4. Scores are assigned conservatively: institutional self-report may support mandate, resources or reported deployment, but performance and outcome points require evidence of exercise, use or observable outputs. Conflicting evidence lowers confidence and may produce a range.

A.2 Missing and conflicting evidence

- Missing public evidence receives no positive score.
- A future deadline or political commitment does not count as technical readiness.
- A deployed infrastructure does not establish access to every intended service.
- An exercise or isolated deployment supports **U=1**, not repeated-outcome evidence.
- A legal framework is not forced into the service score where dimensions are inapplicable.
- Summary bands are descriptive and cannot be compared as if they were interval measurements.

A.3 Sensitivity check

The central proposition remains unchanged under two alternative specifications:

1. removing all institutionally self-reported operationality claims unless independently corroborated lowers the AI Factory and EU-CyCLONe profiles but leaves inherited mechanisms more developed than Plan-specific measures;
2. omitting summary bands entirely and comparing only dimension profiles produces the same directional judgment.

Appendix B. Public Evidence Matrix

ID	Observable proposition	Principal source and evidence stage	Author inference	Counter-evidence or uncertainty	Confidence
E1	Plan published 7 July 2026 with evaluation, access and testing measures	COM(2026) 577; design evidence	identifies Plan-specific functions	implementation not established by the communication	high
E2	cut-off occurred ten days after publication	publication and cut-off dates	outcome attribution is not credible	internal activity may have started	high
E3	CSIRTs, EU-CyCLONe and Reserve pre-dated the Plan	NIS/NIS2, Cyber Solidarity Act, ENISA; design and reported implementation	inherited operational baseline exists	uneven national capacity and limited public performance data	high
E4	Reserve had funding, procedures, exercise and one reported deployment	ENISA; institutional implementation evidence	stronger operational evidence than most Plan-specific measures	depth of specialist coverage not public	moderate-high
E5	European model-evaluation capacity was expected in 2027	Commission governance page; design/implementation planning	service remained prospective	procurement work pre-dated the Plan	high on timing, moderate on readiness

ID	Observable proposition	Principal source and evidence stage	Author inference	Counter-evidence or uncertainty	Confidence
E6	access blueprint and ENISA-JRC platform were future deliverables	Action Plan; design evidence	no accessible service at cut-off	non-public preparatory work possible	high
E7	official pages used different descriptions of AI Factory operationality	Commission pages; institutional self-report	network status must be separated from service status	terminology may reflect different units of analysis	high on divergence, moderate on meaning
E8	official Gigafactory pages report divergent counts of 76 and 77	Commission pages; dynamic institutional evidence	count is source-sensitive but immaterial to central proposition	relative timing and reason for the difference are not public	high
E9	AISI controlled ranges showed constrained multi-step autonomy	AISI technical benchmarks; government evaluation	supports preparedness, not general real-world autonomy	small, vulnerable, undefended environments	moderate
E10	AISI pages report 3/10 and corrected 2/10 for GPT-5.5 on TLO	AISI pages; dynamic technical evidence	benchmark figures require version control	correction concerns grading, not overall direction	high on divergence
E11	finance, electricity and healthcare differ in sector maturity and constraints	ENISA NIS360; agency assessment	absorption capacity will be uneven	ecosystem-level data do not describe every operator	moderate-high
E12	cloud procurement used graduated sovereignty criteria	Commission procurement pages; institutional implementation evidence	dependence is functional, not binary	continuity and substitution were not publicly tested	moderate
E13	Plan-specific measures cluster in design and early mobilisation	synthesis of E1–E12	immediate value is orchestration	rapid future implementation could shorten the delay	moderate

Appendix C. Source Version Note

This note records divergences identified during the version 1.1 integrity check on 18 July 2026. It does not move the substantive research cut-off beyond 17 July 2026.

1. **AI Factories.** The dedicated Commission page stated that nineteen AI Factories and thirteen Antennas were operational and open to European users. A separate Commission topic page stated that work had begun on nineteen factories and that the majority

were expected to be operational by the end of 2026. Version 1.1 therefore distinguishes network designation and general access from site- and service-level operationality.³⁹

2. **AI Gigafactories.** A Commission update of 9 April 2026 referred to 76 responses across 60 sites and 16 Member States. The dedicated project page reported 77 proposals across the same number of sites and Member States. Because the project page does not state a publication or update date, version 1.1 treats 77 as an alternative current official figure rather than necessarily the later figure.⁴⁰
3. **GPT-5.5 cyber range result.** AISI's 13 May 2026 update reported three successful *The Last Ones* runs out of ten. Its dedicated GPT-5.5 evaluation page reports two of ten after a grading correction. Version 1.1 records both figures and avoids treating either as a precise strategic capability estimate.⁴¹

About the Author

Andrea Paone is an independent researcher and technology professional with experience in information technology and systems integration. He holds a master's degree in Strategic Sciences. His work focuses on strategic intelligence, cybersecurity, critical infrastructure, hybrid threats, artificial intelligence, emerging and dual-use technologies, space, biotechnology and international security.

Bibliography

European Commission. *Action Plan on Cybersecurity and Artificial Intelligence*. COM(2026) 577 final. 7 July 2026. CELEX 52026DC0577. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>.

European Commission. "AI Factories." Last updated 23 April 2026. Accessed 18 July 2026. <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>.

³⁹ European Commission, "AI Factories," last updated 23 April 2026, accessed 18 July 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>; European Commission, "Artificial Intelligence," last updated 3 June 2026, accessed 18 July 2026, https://commission.europa.eu/topics/artificial-intelligence_en.

⁴⁰ European Commission, "Commission Marks One Year of the AI Continent Action Plan with Two New Reports on AI Adoption and Policymaking," 9 April 2026, accessed 18 July 2026, <https://digital-strategy.ec.europa.eu/en/news/commission-marks-one-year-ai-continent-action-plan-two-new-reports-ai-adoption-and-policymaking>; European Commission, "AI Gigafactories," accessed 18 July 2026, https://commission.europa.eu/topics/competitiveness/competitiveness-coordination-tool-projects/ai-gigafactories_en.

⁴¹ UK AI Security Institute, "How Fast Is Autonomous AI Cyber Capability Advancing?," 13 May 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>; UK AI Security Institute, "Our Evaluation of OpenAI's GPT-5.5 Cyber Capabilities," 30 April 2026, subsequently updated with a grading correction, accessed 18 July 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-openais-gpt-5-5-cyber-capabilities>; UK AI Security Institute, "Our Evaluation of Claude Mythos Preview's Cyber Capabilities," 13 April 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-claude-mythos-previews-cyber-capabilities>.

European Commission. “AI Gigafactories.” Accessed 18 July 2026. https://commission.europa.eu/topics/competitiveness/competitiveness-coordination-tool-projects/ai-gigafactories_en.

European Commission. “Commission Advances Cloud Sovereignty through Strategic Procurement.” 17 April 2026. https://commission.europa.eu/news-and-media/news/commission-advances-cloud-sovereignty-through-strategic-procurement-2026-04-17_en.

European Commission. “EU AI Office Launches €9 Million Tender for Technical Support on GPAI Safety.” 10 July 2025. <https://digital-strategy.ec.europa.eu/en/funding/eu-ai-office-launches-eu9-million-tender-technical-support-gpai-safety>.

European Commission. “Governance and Enforcement of the AI Act.” Last updated 7 July 2026. <https://digital-strategy.ec.europa.eu/en/policies/ai-act-governance-and-enforcement>.

European Commission. “Sovereign Cloud Framework Explained.” 1 June 2026. https://commission.europa.eu/news-and-media/news/sovereign-cloud-framework-explained-2026-06-01_en.

European Commission. “Strengthening Europe’s Tech Sovereignty.” 3 June 2026. https://commission.europa.eu/news-and-media/news/strengthening-europes-tech-sovereignty-2026-06-03_en.

European Commission. “Artificial Intelligence.” Last updated 3 June 2026. Accessed 18 July 2026. https://commission.europa.eu/topics/artificial-intelligence_en.

European Commission. “Commission Marks One Year of the AI Continent Action Plan with Two New Reports on AI Adoption and Policymaking.” 9 April 2026. Accessed 18 July 2026. <https://digital-strategy.ec.europa.eu/en/news/commission-marks-one-year-ai-continent-action-plan-two-new-reports-ai-adoption-and-policymaking>.

European Parliament and Council of the European Union. Directive (EU) 2016/1148 of 6 July 2016 concerning measures for a high common level of security of network and information systems across the Union. OJ L 194, 19 July 2016. ELI: <https://data.europa.eu/eli/dir/2016/1148/oj>.

European Parliament and Council of the European Union. Directive (EU) 2022/2555 of 14 December 2022 on measures for a high common level of cybersecurity across the Union (NIS2). OJ L 333, 27 December 2022. ELI: <https://data.europa.eu/eli/dir/2022/2555/oj>.

European Parliament and Council of the European Union. Directive (EU) 2022/2557 of 14 December 2022 on the resilience of critical entities. OJ L 333, 27 December 2022. ELI: <https://data.europa.eu/eli/dir/2022/2557/oj>.

European Parliament and Council of the European Union. Regulation (EU) 2022/2554 of 14 December 2022 on digital operational resilience for the financial sector. OJ L 333, 27 December 2022. ELI: <https://data.europa.eu/eli/reg/2022/2554/oj>.

European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). OJ L, 2024/1689, 12 July 2024. ELI: <https://data.europa.eu/eli/reg/2024/1689/oj>.

European Parliament and Council of the European Union. Regulation (EU) 2024/2847 of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements

(Cyber Resilience Act). OJ L, 2024/2847, 20 November 2024. ELI: <https://data.europa.eu/eli/reg/2024/2847/oj>.

European Parliament and Council of the European Union. Regulation (EU) 2025/38 of 19 December 2024 laying down measures to strengthen solidarity and capacities in the Union to detect, prepare for and respond to cyber threats and incidents (Cyber Solidarity Act). OJ L, 2025/38, 15 January 2025. ELI: <https://data.europa.eu/eli/reg/2025/38/oj>.

European Union Agency for Cybersecurity (ENISA). *ENISA NIS360: Latest Insights in the Cybersecurity Maturity and Criticality of NIS Sectors of High Criticality*. May 2026. <https://doi.org/10.2824/8928447>.

European Union Agency for Cybersecurity (ENISA). *ENISA Threat Landscape 2025*. Version 1.2. 9 January 2026. <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>.

European Union Agency for Cybersecurity (ENISA). “EU Cybersecurity Reserve.” Accessed 17 July 2026. <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cybersecurity-reserve>.

European Union Agency for Cybersecurity (ENISA). “EU CyCLONe.” Accessed 17 July 2026. <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cyclone>.

European Union Agency for Cybersecurity (ENISA). “The EU Cybersecurity Reserve - FAQ.” Accessed 17 July 2026. <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cybersecurity-reserve/the-eu-cybersecurity-reserve-faq>.

European Union Agency for Cybersecurity (ENISA). “Cyber Europe 2026: All Eyes on the EU’s Collective Response and Resilience.” 11 June 2026. <https://www.enisa.europa.eu/news/cyber-europe-2026-all-eyes-on-the-eus-collective-response-and-resilience>.

European Union Agency for Cybersecurity (ENISA). *NIS Investments 2025*. 8 December 2025. <https://doi.org/10.2824/7442427>.

European Union Agency for Cybersecurity (ENISA). “Single Reporting Platform (SRP).” Accessed 17 July 2026. <https://www.enisa.europa.eu/topics/product-security-and-certification/single-reporting-platform-srp>.

Happe, Andreas, Aaron Kaplan, and Jürgen Cito. “LLMs as Hackers: Autonomous Linux Privilege Escalation Attacks.” *Empirical Software Engineering* 31, article 70 (2026). <https://doi.org/10.1007/s10664-025-10758-3>.

Tamuka, Nyashadzashe, Topside Ehleketani Mathonsi, Olwal Thomas Otieno, Solly Maswikaneng, Tonderai Muchenje, and Tshimangadzo Mavin Tshilongamulenzhe. “Securing LLM-Based Agents against Cyberattacks: A Comprehensive Survey on Attack Techniques and Defense Strategies.” *Journal of Computer Virology and Hacking Techniques* 22, article 38 (2026). <https://doi.org/10.1007/s11416-026-00622-3>.

UK AI Security Institute. “Our Evaluation of OpenAI’s GPT-5.5 Cyber Capabilities.” 30 April 2026, subsequently updated with a grading correction. Accessed 18 July 2026. <https://www.aisi.gov.uk/blog/our-evaluation-of-openais-gpt-5-5-cyber-capabilities>.

UK AI Security Institute. “How Fast Is Autonomous AI Cyber Capability Advancing?” 13 May 2026. <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>.

UK AI Security Institute. “Measuring AI Agents’ Progress on Multi-Step Cyber Attack Scenarios.” Government technical benchmark, 16 March 2026. <https://www.aisi.gov.uk/research/measuring-ai-agents-progress-on-multi-step-cyber-attack-scenarios>.

UK AI Security Institute. “Our Evaluation of Claude Mythos Preview’s Cyber Capabilities.” 13 April 2026. <https://www.aisi.gov.uk/blog/our-evaluation-of-claude-mythos-previews-cyber-capabilities>.
