



Aletheia Technologies



REPORT · RAPID STRATEGIC ASSESSMENT NO. 01

From Rules to Resilience

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

Andrea Paone

Aletheia Technologies · 2026 · 5 August 2026

AT-RR-2026-001 · Version 2.0 · 2026-08-05

DOI 10.5281/zenodo.21810205

Record: aletheia-technologies.it/verify/r/AVR-AT-RR-2026-001-v2.0-J3IOM/



0 of 5

What an evaluation programme needs and the public record establishes as delivered: personnel, protected compute, published methods, an independent evaluator structure, protected environments. Protected compute scores partial rather than absent; the other four are unevicenced.

This assessment, section 1; Regulation (EU) 2024/1689, Article 113

7 of 8

Key Actions in the Plan that carry a target date. The eighth, on protected compute, carries none.

COM(2026) 577 final; read 5 August 2026

3

Divergences between official sources still live at the research cut-off, reported rather than resolved.

This assessment, Appendix B

Contents

Publication Information	3
Abstract	4
Executive Summary	5
Key Judgments	6
1 Introduction and Research Question	7
2 Scope, Method and Search Boundaries	8
2.1 Product category	8
2.2 Evidence stages	8
2.3 Confidence	9
2.4 Search boundary	9
2.5 Additionality rules	10
2.6 Readiness dimensions	10
2.7 Dynamic sources	10
3 AI-Enabled Cyber Threats	10
3.1 What the evidence supports	10
3.2 Multi-step capability, and how it is scaling	11
3.3 New attack surfaces	13
3.4 Calibration	13
4 What the Union Already Possessed	13
5 Legal Authority and Model Access after 2 August 2026	15
6 What the Action Plan Sequences	17
6.1 The Plan’s own schedule	17
6.2 European model-evaluation capacity	18
6.3 Structured access to advanced systems	18
6.4 ENISA–JRC secure testing platform	18
6.5 Defensive AI, vulnerability management and open source	19
6.6 Grand Challenge	20

6.7	AI Factories, Gigafactories and protected compute	20
6.8	Reading the measures together	21
7	Critical Infrastructure: Uneven Absorption	21
8	Strategic Dependencies	22
9	Alternative Interpretations and Falsifiers	24
9.1	Normal early-stage implementation	24
9.2	Regulatory layering	24
9.3	Rapid operational convergence	25
9.4	Capability growth with continued external dependence	25
9.5	Falsifiers of the historical assessment	25
10	Indicator Register	25
11	Conclusion	27
	Limitations	28
	Appendix A. Measurement Rules	28
	Appendix B. Source Divergences	29
	Appendix C. Five-Dimension Readiness Profile	29
	Bibliography	31
11.1	Union legislation and official documents	31
11.2	European Commission web sources	31
11.3	European Union Agency for Cybersecurity	32
11.4	Technical benchmarks and academic literature	33
	About the Author	33

Publication Information

Author: Andrea Paone

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

Product category: Rapid Strategic Assessment

Version: 2.0, 5 August 2026

Preceding versions: 1.0 (17 July 2026), 1.1 (18 July 2026)

Research cut-off: 5 August 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/>

Version DOI: 10.5281/zenodo.21810205

Concept DOI: 10.5281/zenodo.21417540

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

Editorial imprint: Aletheia Technologies — Strategic Intelligence & Research

Status: Independent Research Initiative; not a legally incorporated entity

Review status: Independent Rapid Strategic Assessment; 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 2.0 (5 August 2026), DOI: 10.5281/zen-

odo.21810205.

What changed in version 2.0. The research cut-off moves from 17 July to 5 August 2026, and the central proposition changes: version 1.1 described the Plan as an orchestration and sequencing framework, which restated the Plan's own timetable; this version identifies the conversion problem between legal authority and operational capacity, and treats the Plan's sequencing as a response to it. The enforcement position under the AI Act changed on 2 August 2026 and is analysed before the Plan's own measures, because the problem must be visible before the response to it. Eight Key Judgments become six.

Four corrections follow from reading primary sources rather than summaries: the Plan does not mention the CSIRTs Network or EU-CyCLONe, so those are no longer attributed to it; the open-source resilience measure is reclassified as Plan-specific; a July 2026 ENISA publication that fell inside the previous cut-off is now cited, which reverses a negative finding while establishing that it cannot evidence Plan delivery, since it was published on the day of adoption; and two ENISA reports previously cited unread have been read.

Three constructs are withdrawn as insufficiently supported: the R0–R4 summary bands, the fifteen-cell sector matrix and the twenty-five-cell dependency matrix. The five-dimension readiness profile moves to Appendix C as a diagnostic. Advanced-accelerator dependence is recorded as an evidence gap rather than a finding. Full measurement detail is in the supplementary evidence package.

Supplementary evidence package. Every table, rating and negative-evidence statement in this paper maps to a row in the accompanying package: claim ledger, source audit register, measurement and scoring register, sector and dependency registers, indicator register, sensitivity tests and negative-evidence register.

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 asks whether the European Union can convert a regulator's legal power to obtain access to advanced AI models into a usable capability to evaluate them, and to defend critical infrastructure against AI-enabled cyber threats. It is an ex-ante assessment of additionality, implementation readiness and capability-development pathways, not an impact evaluation. The research cut-off is 5 August 2026.

The method separates what the Union inherited from what the 2026 Action Plan on Cybersecurity and Artificial Intelligence introduced, and assigns each proposition to an evidence stage: policy design, institutionally reported implementation, independently corroborated implementation, or observed performance. The implementation update rests on a targeted scan of official and institutional sources rather than an exhaustive search of Member State and commercial records, and every negative finding is stated within that boundary.

The central finding is a conversion problem, stated in full in section 1. Union law had the access power and the published procedure; the reviewed public record had no completed evaluation, no allocated protected compute, no published methodology and no repeatedly usable service. The Action Plan's contribution is to sequence the missing functions around a baseline it inherited.

Two further findings frame the problem. Independent benchmark evidence shows autonomous cyber capability improving with inference-time compute and across model generations in bounded, undefended ranges, without demonstrating reliable operation against

defended real-world industrial systems. And sectoral capacity to absorb common European capabilities is likely to remain uneven, on ecosystem-level evidence that does not describe individual operators.

Contribution: descriptive, comparative and analytical.

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

Executive Summary

On 2 August 2026 something changed that the Action Plan did not cause. The Commission's supervisory powers over providers of general-purpose AI models became applicable, and with them a power to require access to a model — through interfaces, other technical means, or source code — for the purpose of conducting an evaluation, with refusal exposed to fines of up to 3 per cent of worldwide turnover. A Union authority now holds an enforceable claim on the systems it is expected to assess.

Authority is not capability. The reviewed public record does not establish the personnel, the compute, the published methods, the independent evaluator structure or the protected environments that an evaluation programme requires. The procedural side is further advanced than the operational side. The Commission adopted the detailed arrangements for these proceedings on 20 July 2026 and published them the following day, but they were not due to take effect until 10 August, five days after this cut-off.¹ The Commission's own governance page still places European evaluation capacity in 2027. The gap between a legal power and a working capability is the subject of this assessment.

The Union did not enter this position empty-handed. It already possessed mechanisms for technical cooperation, crisis coordination, incident response, sectoral supervision, product-security reporting and AI computing, established by the NIS Directive and NIS2, the Cyber Solidarity Act, DORA, the Critical Entities Resilience Directive, the Cyber Resilience Act, the AI Act and EuroHPC. The Action Plan of 7 July 2026 inherited that architecture rather than creating it.²

What the Plan contributes is a sequence for the missing functions. Its eight Key Actions distribute evaluation capacity, structured access, secure testing, vulnerability management, open-source resilience, protected compute and industrial development across target quarters running from the third quarter of 2026 to 2027; the eighth, on protected compute, carries no date. Read against those dates, the absence of services in early August is unremarkable. It is also uninformative: an unexpired deadline is evidence neither of progress nor of failure, and no interim milestone is public against which either could be measured.

¹Commission Implementing Regulation (EU) 2026/1755 of 20 July 2026 on detailed arrangements for the conduct of certain proceedings by the Commission pursuant to Regulation (EU) 2024/1689 of the European Parliament and of the Council, OJ L, 2026/1755, 21 July 2026, CELEX 32026R1755, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1755>. Adopted on the basis of Articles 92(6) and 101(6) of Regulation (EU) 2024/1689. Recitals for the enumeration of access modalities, including application programming interfaces, internal access, source code, model weights, the infrastructure used for hosting the model, and access to inspect and modify the system state during interaction, extending to all levels of access granted to the provider's own employees; provisions on the selection and involvement of independent experts and on the procedural guarantees attaching to Article 101 penalties; Article 15 for entry into force on the twentieth day following publication, that is 10 August 2026.

²European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

Two further findings frame the risk. Autonomous cyber capability is improving on two axes at once — the compute spent at inference time, and the model generation — in ranges that AISI itself describes as small and undefended, with initial access already granted; nothing in that evidence demonstrates reliable operation against defended industrial systems. And sectoral capacity to absorb whatever the Union builds is likely to remain uneven: finance is comparatively well positioned, electricity is institutionally strong but technically specialised, and healthcare combines high criticality with fragmented capacity.

Overall assessment. At 5 August 2026 Union law had moved ahead of publicly evidenced European capacity to use it, and the Plan’s function was to sequence the functions that would close that distance. This is held with moderate confidence: the negative half of the finding rests on a targeted rather than exhaustive search, and public silence is not the same as institutional absence.

Key Judgments

1. **Legal authority and procedure moved ahead of publicly evidenced operational capacity.** From 2 August 2026 the Commission may require access to a general-purpose AI model for the purpose of regulatory evaluation, with non-compliance exposed to fines under Article 101, and the detailed procedural arrangements were published on 21 July 2026 with effect from 10 August.³ The reviewed public record does not establish completed evaluations, allocated protected compute, a published evaluation methodology or repeatedly usable European service capacity. The legal and procedural position is stated with high confidence, since it is read from the acts themselves. The assessment of operational capacity carries only moderate confidence, because it rests on the absence of evidence within a bounded search.
2. **The Plan sequences AI-specific functions around an inherited cyber baseline.** It did not create the Union’s horizontal cooperation, crisis-response, financial-resilience, product-security or compute architecture. Its principal additionality lies in sequencing model evaluation, structured access, secure testing, vulnerability remediation, open-source resilience and protected compute around that pre-existing base. Provenance and target dates are established with high confidence from the instruments themselves. That sequencing is the Plan’s *principal* additionality is an interpretation, and is advanced with moderate confidence.
3. **No Plan-attributable service use or outcome was identified in the reviewed public record.** No qualifying public evidence of Plan-attributable service use, completed evaluation, operational access or outcome was found in the official and institutional sources examined through 5 August 2026. The short interval since adoption and the later target dates make an outcome assessment premature. This judgment is held with moderate confidence, and only within the boundary of the search: official and institutional sources were examined, not an exhaustive record.

³Commission Implementing Regulation (EU) 2026/1755 of 20 July 2026 on detailed arrangements for the conduct of certain proceedings by the Commission pursuant to Regulation (EU) 2024/1689 of the European Parliament and of the Council, OJ L, 2026/1755, 21 July 2026, CELEX 32026R1755, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1755>. Adopted on the basis of Articles 92(6) and 101(6) of Regulation (EU) 2024/1689. Recitals for the enumeration of access modalities, including application programming interfaces, internal access, source code, model weights, the infrastructure used for hosting the model, and access to inspect and modify the system state during interaction, extending to all levels of access granted to the provider’s own employees; provisions on the selection and involvement of independent experts and on the procedural guarantees attaching to Article 101 penalties; Article 15 for entry into force on the twentieth day following publication, that is 10 August 2026.

4. **Cyber-agent capability is improving in controlled ranges, not demonstrated against defended critical infrastructure.** Capability improves with inference-time compute and across model generations in the ranges reviewed. Those ranges are bounded and undefended, with initial access already granted, and the evidence does not demonstrate reliable end-to-end operation against defended, heterogeneous real-world critical infrastructure. Moderate confidence: the trend is measured, but it is measured in conditions that do not resemble a defended target.
5. **Sectoral absorption is likely to remain uneven.** ENISA's ecosystem-level evidence supports material differences among finance, electricity and healthcare. The inference that those differences will produce uneven absorption of common European capabilities is plausible but extends beyond the source's unit of analysis, which assesses ecosystems rather than operators. Moderate confidence follows from that gap between what the source measures and what the judgment claims.
6. **Dependency findings remain narrow and functional.** The strongest documented constraints concern access to frontier models and specialist organisational capacity. The cloud evidence supports concern about continuity and substitutability without demonstrating failure under stress. Advanced-accelerator dependence is an evidence gap in this assessment and is not presented as a verified concentration of European dependence. Model access and specialist capacity are assessed with moderate confidence; the broader dependency picture with low to moderate confidence, since parts of it rest on a single procurement and parts on no cited source at all.

Two findings that occupied Key Judgment slots in earlier versions are retained as supporting evidence rather than principal conclusions: that ENISA guidance aligned with Key Action 4 already existed on the day the Plan was adopted, and that the Plan carries eight Key Actions with target dates on seven of them. The rule that each measure must be assessed at the evidence stage appropriate to its object type belongs to the method and is stated in section 2.

1 Introduction and Research Question

By mid-2026 the European Union had assembled an extensive body of law and institutions covering network security, product security, crisis coordination, critical-entity resilience, financial-sector resilience and the governance of advanced AI. The open question was no longer whether rules existed, but whether they could support operational action against a threat environment changing faster than legislative cycles.

Generative and agentic systems already assist vulnerability research, social engineering, code production and reconnaissance, and they equally assist detection, triage and remediation. Which side benefits more depends on who can access, evaluate, integrate and govern the systems.

The Commission published the *Action Plan on Cybersecurity and Artificial Intelligence* on 7 July 2026.⁴ Four weeks later the AI Act's supervisory provisions became applicable. This paper asks four questions across the two instruments:

⁴European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

1. What legal authority does a Union body hold over advanced AI systems for the purpose of evaluating them, and since when?
2. What operational capacity to exercise that authority is publicly evidenced?
3. What does the Action Plan add that is genuinely its own, and in what sequence?
4. Where is the capacity to absorb common European capabilities likely to be uneven, and where does dependence bind?

Central proposition. At the 5 August 2026 cut-off, Union law had established a sanction-backed access power and had published the procedural arrangements for Commission-led evaluations of general-purpose AI models, although those arrangements were not due to take effect until 10 August. The reviewed public record still did not establish completed evaluations, allocated protected compute, a published evaluation methodology or repeatedly usable European service capacity. The Action Plan’s distinctive contribution remained to sequence broader organisational access, secure sectoral testing, remediation and compute functions around a horizontal cyber baseline it inherited rather than created. The proposition is advanced with moderate confidence, for the reason given in section 9.1.

The legal component of that proposition is established directly from the Regulation. The operational component is a negative finding bounded by the search described in section 2.4, and the strongest case against the proposition — that the public record may show a transparency gap rather than a capability gap — is set out in section 9.1.

2 Scope, Method and Search Boundaries

2.1 Product category

This is a Rapid Strategic Assessment: an ex-ante assessment of additionality, implementation readiness and capability-development pathways. It is not an impact assessment, does not measure changes in European cyber performance caused by the Plan, and is not an official evaluation. It is not peer reviewed.

The analysis covers the Union, its institutions and agencies, and national or sectoral implementation where comparable public evidence exists. Finance, electricity and healthcare serve as comparative cases, selected to maximise variation in regulatory maturity, technical architecture and organisational fragmentation. They are illustrative, not statistically representative.

2.2 Evidence stages

Every material proposition is assigned a stage, and no claim is stated above its stage:

1. **design** — legislation, communications, mandates, announced objectives;
2. **institutionally reported implementation** — procurement, staffing, platform development or deployment reported by the responsible body;
3. **independently corroborated implementation** — auditors, users, operators, procurement records, independent technical assessment;
4. **performance or outcome** — repeated use, service availability, measurable effect.

An institution is authoritative about what it states. An announcement, a funding allocation or a self-reported status is not evidence that a service is available, accessible or effective. Where two official pages conflict, both are reported and the permitted inference is narrowed rather than resolved by preference.

Each measure is assessed at the evidence stage appropriate to its object type. A legal framework, a coordination network, a response service, a compute infrastructure, a guidance document and an unbuilt platform are not made comparable by assigning each a single label, and this assessment does not compare them as though they were.

2.3 Confidence

Every assessment in this paper carries a stated level of confidence, expressed in prose rather than as a label. The levels mean the following.

High confidence rests on multiple independent sources, or on a single authoritative primary source for a claim about that source's own content; the evidence is direct and complete, and the assessment does not move under the alternative specifications tested in Appendix A.

Moderate confidence rests on authoritative but self-reported or single-sourced evidence, or on direct evidence requiring one inferential step, or on an assessment sensitive to a single assumption.

Low confidence rests on indirect or contested evidence, on an unverified premise, or on a source whose unit of analysis differs from the claim being made.

An official source does not by itself produce high confidence. It produces high confidence only about what the institution states. Confidence attaches to a particular assessment and not to the document as a whole, and where an assessment has a legal component and an evidential one, the two are stated separately because they are not equally secure.

2.4 Search boundary

The implementation update for this version rests on a **targeted scan of official and institutional sources**: the Action Plan and the Regulation read in full; the Commission pages for AI Act governance, AI Factories, Gigafactories, cloud procurement and technology sovereignty; the ENISA pages for the Cybersecurity Reserve, EU-CyCLONe, the Single Reporting Platform, the Threat Landscape and the frontier-AI publication; ENISA NIS360 and NIS Investments 2025; and the AISI paper with its associated evaluation posts. Coverage extends to all eight Key Actions, the AI Office evaluation-capacity initiative, Article 92 implementation and implementing acts, ENISA and JRC outputs, protected-compute allocation, and public evidence of users, completed evaluations, exercises and remediation.

It is **not** an exhaustive search of Member State records, national authority publications, commercial disclosures or procurement databases. Every negative finding in this paper is therefore a statement about that reviewed record and is worded accordingly: no qualifying public evidence was identified. It is not a statement that the activity or capacity does not exist. Where a negative finding is load-bearing, its search is logged in the supplementary negative-evidence register.

The surface searched matters as much as the terms used, and the two are not interchangeable. A negative claim about the **existence of a legal instrument** is answerable only against the legal database: an institution's website records what that institution chooses to describe, while the Official Journal records what the law is. The implementing regulation under Article 92(6) illustrates the point, since it was adopted on 20 July and published on 21 July 2026, two weeks before this cut-off, and is discoverable only through the second of those two surfaces.⁵ The negative findings that remain in this paper concern services, users and

⁵Commission Implementing Regulation (EU) 2026/1755 of 20 July 2026 on detailed arrangements for the conduct of certain proceedings by the Commission pursuant to Regulation (EU) 2024/1689 of the European Parliament and of the Council, OJ L, 2026/1755, 21 July 2026, CELEX 32026R1755, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1755>. Adopted on the basis of Articles 92(6) and 101(6) of Regulation (EU)

outputs rather than instruments, and the supplementary register records the surface each was searched against.

2.5 Additionality rules

Measures are classified as **inherited** (operational or developing before 7 July 2026); **redirected** (a pre-existing mechanism given an AI-cyber role without demonstrated change in resources, scope or timing); **accelerated** (public evidence of changed budget, staffing, procurement, timeline or scope); or **new and Plan-specific** (introduced by the Plan). A measure is not called accelerated merely because the Plan mentions it.

2.6 Readiness dimensions

Five dimensions are recorded for operational mechanisms and service-like measures, and are tabulated in Appendix C: mandate and governance (**G**), resources and implementation machinery (**R**), technical readiness (**T**), access and integration (**A**), and operational use and effects (**U**). Each takes one of:

notation	meaning
2	established evidence meeting the criterion
1	partial, planned, pilot or limited evidence
0	affirmative evidence that the criterion is not met
NE	no qualifying public evidence identified
NA	not applicable to this object type
C	conflicting evidence

The distinction between **0** and **NE** is deliberate. Version 1.1 used a single zero for both, which allowed an absence of published evidence to read as an absence of capability.

Earlier versions summed these dimensions into bands labelled R0 to R4. Those bands are withdrawn, for the reason given in Appendix C, and the profile itself is a supplementary diagnostic reported there rather than an instrument of the argument.

2.7 Dynamic sources

Claims resting on modifiable pages carry an access date, the stated update date where available, the proposition observed, and a checksummed snapshot. Divergences between official pages are reported. Three are live at this cut-off and are listed in Appendix B.

3 AI-Enabled Cyber Threats

3.1 What the evidence supports

Generative models produce persuasive multilingual content, personalise messages and sustain iterative interaction. ENISA identifies AI as an optimisation tool for phishing and social

2024/1689. Recitals for the enumeration of access modalities, including application programming interfaces, internal access, source code, model weights, the infrastructure used for hosting the model, and access to inspect and modify the system state during interaction, extending to all levels of access granted to the provider's own employees; provisions on the selection and involvement of independent experts and on the procedural guarantees attaching to Article 101 penalties; Article 15 for entry into force on the twentieth day following publication, that is 10 August 2026.

engineering, and as an emerging target through attacks on AI systems and supply chains.⁶ The change is one of scale, speed and accessibility rather than of kind: an attacker still needs target data, delivery infrastructure, credential collection, access and a method of exploitation.

Peer-reviewed experiments on autonomous Linux privilege escalation show model-based agents succeeding in bounded controlled scenarios while failing on command validity, common-sense reasoning, error handling and multi-step exploitation with temporal dependencies.⁷ That is meaningful component capability, not reliable operation against unknown defended networks.

3.2 Multi-step capability, and how it is scaling

The strongest public evidence comes from the UK AI Security Institute's purpose-built cyber ranges: *The Last Ones*, a 32-step corporate network attack, and *Cooling Tower*, a 7-step industrial control system attack.⁸ Two trends matter more than any single score.

Capability scales with compute spent at inference time. Across seven models released between August 2024 and February 2026, performance rises log-linearly with token spend, with no plateau observed up to 100 million tokens; moving from 10 million to 100 million tokens yields gains of up to 59 per cent. AISI states that this requires no specific technical sophistication from the operator.⁹

Capability rises across model generations at fixed budgets. On the corporate range at a 10-million-token budget, average steps completed rose from 1.7 for a model released in August 2024 to 9.8 for one released in February 2026.¹⁰

Against those trends, absolute performance remains bounded. In April 2026 Claude Mythos Preview became the first model to complete *The Last Ones* end to end, in three of ten attempts, completing an average of 22 of 32 steps across all attempts; the next best model averaged 16.¹¹ A May 2026 update reported a newer checkpoint completing *The Last Ones* in

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

⁷Andreas 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>.

⁸Linus Folkerts et al., "Measuring AI Agents' Progress on Multi-Step Cyber Attack Scenarios," arXiv:2603.11214v3, submitted 11 March 2026, current version 17 March 2026, <https://arxiv.org/abs/2603.11214>. Abstract and section 2 for the log-linear scaling of performance with inference-time compute with no observed plateau to 100 million tokens, the gain of up to 59 per cent between 10 and 100 million tokens, the generational rise from 1.7 to 9.8 average steps at a 10-million-token budget, the best single run of 22 of 32 steps against an estimated 14 human hours, and the industrial-range average of 1.2 to 1.4 steps of 7 with a maximum of 3; section 2 for the description of both ranges as sequential attack chains without active defenders or detection mechanisms.

⁹Linus Folkerts et al., "Measuring AI Agents' Progress on Multi-Step Cyber Attack Scenarios," arXiv:2603.11214v3, submitted 11 March 2026, current version 17 March 2026, <https://arxiv.org/abs/2603.11214>. Abstract and section 2 for the log-linear scaling of performance with inference-time compute with no observed plateau to 100 million tokens, the gain of up to 59 per cent between 10 and 100 million tokens, the generational rise from 1.7 to 9.8 average steps at a 10-million-token budget, the best single run of 22 of 32 steps against an estimated 14 human hours, and the industrial-range average of 1.2 to 1.4 steps of 7 with a maximum of 3; section 2 for the description of both ranges as sequential attack chains without active defenders or detection mechanisms.

¹⁰Linus Folkerts et al., "Measuring AI Agents' Progress on Multi-Step Cyber Attack Scenarios," arXiv:2603.11214v3, submitted 11 March 2026, current version 17 March 2026, <https://arxiv.org/abs/2603.11214>. Abstract and section 2 for the log-linear scaling of performance with inference-time compute with no observed plateau to 100 million tokens, the gain of up to 59 per cent between 10 and 100 million tokens, the generational rise from 1.7 to 9.8 average steps at a 10-million-token budget, the best single run of 22 of 32 steps against an estimated 14 human hours, and the industrial-range average of 1.2 to 1.4 steps of 7 with a maximum of 3; section 2 for the description of both ranges as sequential attack chains without active defenders or detection mechanisms.

¹¹UK AI Security Institute, "Our Evaluation of Claude Mythos Preview's Cyber Capabilities," 13 April 2026, accessed 5 August 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-claude-mythos-previews-cyber-capabilities>, for the 32-step description of *The Last Ones*, the 20-hour human estimate, the first end-to-end completion in three of ten attempts, the average of 22 of 32 steps across attempts, and the next best model's average of 16 steps.

six of ten attempts and *Cooling Tower*, previously unsolved, in three of ten.¹² On the industrial range the paper’s population averaged 1.2 to 1.4 steps of 7, with a maximum of 3.¹³

The two figures for the same range are not in conflict: the paper covers models to February 2026, and the checkpoint that solved *Cooling Tower* was evaluated later. Reporting them together, with their populations, is more informative than reconciling them.

What the ranges do not test. AISI describes them as small, undefended enterprise networks where initial access has already been gained.¹⁴ Three limitations follow, stated by AISI itself. The two ranges “lack the active defenders, defensive tooling, and alert penalties that real-world environments typically have”. The cyber tasks test skills in isolation. And no conclusion follows about a well-defended target: testing is scoped to directed attacks on vulnerable targets where network access already exists. Further ranges addressing detection evasion on hardened targets are under construction.¹⁵

One detail sharpens the industrial picture. A model’s failure on the industrial range occurred on its IT sections rather than its OT-specific steps, so that result says nothing about industrial-control competence.¹⁶

Operational technology compounds the difficulty: proprietary protocols, long-lived equipment, safety constraints, segmented networks and operator-specific processes. Converting access into a physical effect normally requires domain knowledge, persistence and interaction with the target environment.

¹²UK AI Security Institute, “How Fast Is Autonomous AI Cyber Capability Advancing?,” 13 May 2026, accessed 5 August 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>, for the characterisation of the ranges as small, undefended enterprise networks where initial access has already been gained, and for the newer checkpoint completing *The Last Ones* in six of ten attempts and *Cooling Tower* in three of ten.

¹³Linus Folkerts et al., “Measuring AI Agents’ Progress on Multi-Step Cyber Attack Scenarios,” arXiv:2603.11214v3, submitted 11 March 2026, current version 17 March 2026, <https://arxiv.org/abs/2603.11214>. Abstract and section 2 for the log-linear scaling of performance with inference-time compute with no observed plateau to 100 million tokens, the gain of up to 59 per cent between 10 and 100 million tokens, the generational rise from 1.7 to 9.8 average steps at a 10-million-token budget, the best single run of 22 of 32 steps against an estimated 14 human hours, and the industrial-range average of 1.2 to 1.4 steps of 7 with a maximum of 3; section 2 for the description of both ranges as sequential attack chains without active defenders or detection mechanisms.

¹⁴UK AI Security Institute, “How Fast Is Autonomous AI Cyber Capability Advancing?,” 13 May 2026, accessed 5 August 2026, <https://www.aisi.gov.uk/blog/how-fast-is-autonomous-ai-cyber-capability-advancing>, for the characterisation of the ranges as small, undefended enterprise networks where initial access has already been gained, and for the newer checkpoint completing *The Last Ones* in six of ten attempts and *Cooling Tower* in three of ten.

¹⁵UK AI Security Institute, “Our Evaluation of OpenAI’s GPT-5.5 Cyber Capabilities,” 30 April 2026, subsequently updated with a grading correction, accessed 5 August 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-openai-gpt-5-5-cyber-capabilities>, for the statement that the two ranges lack the active defenders, defensive tooling and alert penalties that real-world environments typically have, that cyber tasks test skills in isolation, that no conclusion follows about a well-defended target, that testing is scoped to directed attacks on vulnerable targets where network access already exists, that further ranges addressing detection evasion on hardened targets are under construction, and that one model’s industrial-range failure occurred on the IT sections rather than the OT-specific steps.

¹⁶UK AI Security Institute, “Our Evaluation of OpenAI’s GPT-5.5 Cyber Capabilities,” 30 April 2026, subsequently updated with a grading correction, accessed 5 August 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-openai-gpt-5-5-cyber-capabilities>, for the statement that the two ranges lack the active defenders, defensive tooling and alert penalties that real-world environments typically have, that cyber tasks test skills in isolation, that no conclusion follows about a well-defended target, that testing is scoped to directed attacks on vulnerable targets where network access already exists, that further ranges addressing detection evasion on hardened targets are under construction, and that one model’s industrial-range failure occurred on the IT sections rather than the OT-specific steps.

3.3 New attack surfaces

Agentic systems combine planning, memory and tool use, often with access to logs, cloud controls or remediation systems. They are exposed to prompt manipulation, memory or retrieval poisoning, malicious tool-output injection, unsafe tool use, model manipulation and supply-chain compromise.¹⁷ A defensive agent is therefore a privileged target, and model evaluation alone is insufficient: security must cover integrations, permissions, data sources, monitoring, human oversight and fail-safe behaviour.

3.4 Calibration

The progression runs from AI-assisted content and analysis, through automation of bounded tasks and multi-step agentic assistance, to constrained autonomy in defined environments and finally reliable end-to-end autonomy in unknown defended environments. Public evidence strongly supports the first three levels. The fourth is now demonstrated in controlled ranges, inconsistently and with the limitations AISI states. The fifth is not demonstrated for defended critical infrastructure in the reviewed public evidence.

The compute-scaling result adds a variable rather than replacing the others. Within the range studied, capability rises with token spend and AISI states this requires no particular operator sophistication, so how much an adversary is willing to spend per operation becomes one input into a threat estimate. It does not displace the model used, the agent scaffolding, tool access, the target environment, specialist domain knowledge, detection, active defences, or the reliability and variance of the attempt. The scaling was also measured on undefended ranges, so extrapolating it to defended environments is not licensed by the evidence. Europe therefore has reason to build evaluation and preparedness capacity, and no basis for presenting controlled benchmark success as proof of real-world autonomy.

4 What the Union Already Possessed

The NIS Directive established the CSIRTs Network.¹⁸ NIS2 strengthened it and formalised EU-CyCLONe, launched in 2020, with ENISA supporting both.¹⁹ ²⁰ These provide routes for sharing and coordinating information on AI-enabled threats, although national capability and legal implementation remain uneven. Neither body is mentioned in the Action Plan; their relevance to the AI-cyber agenda is established instead by ENISA, which records their input to its July 2026 recommendations.²¹

¹⁷Nyashadzashe Tamuka, Onderai 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* (2026), <https://doi.org/10.1007/s11416-026-00622-3>.

¹⁸Directive (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, OJ L 194, 19 July 2016, ELI: <https://data.europa.eu/eli/dir/2016/1148/oj>.

¹⁹Directive (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, OJ L 333, 27 December 2022, ELI: <https://data.europa.eu/eli/dir/2022/2555/oj>.

²⁰European Union Agency for Cybersecurity, “EU CyCLONe,” accessed 5 August 2026, <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cyclone>.

²¹European Union Agency for Cybersecurity, *ENISA’s View on Cybersecurity in the Frontier AI Era*, TLP:CLEAR, July 2026, https://www.enisa.europa.eu/sites/default/files/2026-07/ENISA%20view%20on%20cybersecurity%20in%20the%20frontier%20AI%20era_en_0.pdf, at the acknowledgements for CSIRTs Network and EU-CyCLONe input, and in the recommendations for Union-wide security-evaluation benchmarks including standardised cyber-range testing, exploitability metrics and chained-attack simulations. The document states the month of publication and not the day.

The Cyber Solidarity Act created the European Cybersecurity Alert System, a Cyber Emergency Mechanism and the European Cybersecurity Reserve.²² By mid-2026 the Reserve had moved beyond a legislative commitment. ENISA administers it under a Contribution Agreement signed with the Commission in August 2025, funded by EUR 36 million earmarked from the Digital Europe Programme over three years; trusted providers are procured through open procedure and the list is being renewed; requests reach ENISA through NIS2 single points of contact and CERT-EU.²³ ²⁴ It was first deployed to Moldova in September 2025, and in June 2026 it was tested for the first time under Cyber Europe, activated through a scenario in which participants followed ENISA's standard operating procedure.²⁵ ²⁶ That is one deployment and one exercise: real operational evidence, and not yet evidence of repeated use.

DORA has applied since January 2025 and gives the financial sector a structured regime for ICT risk, incident reporting, resilience testing and third-party oversight.²⁷ It creates no AI-specific defence but supplies processes into which AI-enabled risk can be integrated. The Critical Entities Resilience Directive adds the physical-resilience layer.²⁸ The Cyber Resilience Act created product-security and vulnerability-handling obligations; its Single Reporting Platform was being prepared for mandatory reporting from 11 September 2026, and the Act becomes fully applicable by the end of 2027.²⁹ ³⁰ ³¹

The AI Act requires providers of general-purpose models with systemic risk to evaluate and adversarially test them, mitigate systemic risks, report serious incidents and maintain

²² 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, OJ L, 2025/38, 15 January 2025, ELI: <https://data.europa.eu/eli/reg/2025/38/oj>.

²³ European Union Agency for Cybersecurity, "EU Cybersecurity Reserve," accessed 5 August 2026, <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cybersecurity-reserve>, for the EUR 36 million Digital Europe Programme allocation.

²⁴ European Union Agency for Cybersecurity, "The EU Cybersecurity Reserve — FAQ," accessed 5 August 2026, <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cybersecurity-reserve/the-eu-cybersecurity-reserve-faq>, for the August 2025 Contribution Agreement, the three-year earmark, provider procurement, the request route through NIS2 single points of contact and CERT-EU, and the first deployment to Moldova in September 2025.

²⁵ European Union Agency for Cybersecurity, "The EU Cybersecurity Reserve — FAQ," accessed 5 August 2026, <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cybersecurity-reserve/the-eu-cybersecurity-reserve-faq>, for the August 2025 Contribution Agreement, the three-year earmark, provider procurement, the request route through NIS2 single points of contact and CERT-EU, and the first deployment to Moldova in September 2025.

²⁶ European Union Agency for Cybersecurity, "Cyber Europe 2026: All Eyes on the EU's Collective Response and Resilience," accessed 5 August 2026, <https://www.enisa.europa.eu/news/cyber-europe-2026-all-eyes-on-the-eus-collective-response-and-resilience>, for the first test of the Reserve under Cyber Europe and its activation through a scenario following ENISA's standard operating procedure.

²⁷ Regulation (EU) 2022/2554 of the European Parliament and of the Council 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>.

²⁸ Directive (EU) 2022/2557 of the European Parliament and of the Council 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>.

²⁹ 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, OJ L, 2024/2847, 20 November 2024, ELI: <https://data.europa.eu/eli/reg/2024/2847/oj>.

³⁰ European Union Agency for Cybersecurity, "Single Reporting Platform (SRP)," accessed 5 August 2026, <https://www.enisa.europa.eu/topics/product-security-and-certification/single-reporting-platform-srp>.

³¹ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

cybersecurity protections.³² Until August 2026 those obligations sat alongside a supervisory apparatus that could not yet be enforced; section 6 returns to what changed.

EuroHPC and the AI Factories form a growing general-purpose compute base. Official Commission descriptions of their status diverge. One page presents nineteen AI Factories and thirteen associated Antennas as operational and open to European users. A separate overview states that work had begun on nineteen, and that the majority were expected to be operational by the end of 2026.³³ ³⁴ The two descriptions are not necessarily irreconcilable, since network designation, partial service availability and site-level operationality 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 industrial-control capability from a general infrastructure label.

The baseline was substantial and heterogeneous: mature cooperation and crisis mechanisms, developing product-security tools, limited AI-specific technical capacity. The Plan inherited it.

5 Legal Authority and Model Access after 2 August 2026

The AI Act entered application on 2 August 2026. Its Article 113 brought forward some provisions to 2 August 2025 — including Chapter V on general-purpose AI models — but expressly excepted Article 101 from that advance, and did not advance Chapter IX, which contains the supervisory powers.³⁵ The consequence is that the substantive obligations of general-purpose AI model providers applied from August 2025, while the Commission’s powers to supervise and enforce them, and the fines attached, applied from **2 August 2026**.

Three of those powers matter here. The AI Office may conduct evaluations of a general-purpose AI model to assess compliance where information gathered under Article 91 is insufficient, or to investigate systemic risks at Union level. For that purpose the Commission “may request access to the general-purpose AI model concerned through APIs or further appropriate technical means and tools, including source code”, stating the legal basis, purpose and period, and the fines applicable for failure to provide access. Failure to make access available for such an evaluation is subject to fines of up to 3 per cent of worldwide annual turnover or EUR 15 million, whichever is higher. The Commission may appoint independent experts to carry out evaluations on its behalf, and Article 92(6) requires the detailed arrangements

³²Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, OJ L, 2024/1689, 12 July 2024, ELI: <https://data.europa.eu/eli/reg/2024/1689/oj>, as amended by Regulation (EU) 2026/1744.[omnibus] Article 113 on entry into force and application; Article 91 on the power to request documentation and information; Article 92 on the power to conduct evaluations, in particular Article 92(3) on access through APIs or further appropriate technical means and tools including source code, Article 92(4) on the content of the request and Article 92(6) on implementing acts; Article 93 on the power to request measures; Article 101(1) on fines for providers of general-purpose AI models, in particular Article 101(1)(d).

³³European Commission, “AI Factories,” accessed 5 August 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>.

³⁴European Commission, “Artificial Intelligence,” accessed 5 August 2026, https://commission.europa.eu/topics/artificial-intelligence_en.

³⁵Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, OJ L, 2024/1689, 12 July 2024, ELI: <https://data.europa.eu/eli/reg/2024/1689/oj>, as amended by Regulation (EU) 2026/1744.[omnibus] Article 113 on entry into force and application; Article 91 on the power to request documentation and information; Article 92 on the power to conduct evaluations, in particular Article 92(3) on access through APIs or further appropriate technical means and tools including source code, Article 92(4) on the content of the request and Article 92(6) on implementing acts; Article 93 on the power to request measures; Article 101(1) on fines for providers of general-purpose AI models, in particular Article 101(1)(d).

to be set out in implementing acts.³⁶ Those arrangements exist. Commission Implementing Regulation (EU) 2026/1755 of 20 July 2026, adopted under Articles 92(6) and 101(6), lays down detailed arrangements for the conduct of these proceedings.³⁷ It specifies that access may include application programming interfaces, internal access, access to source code, access to **model weights**, access to the infrastructure used for hosting the model, and access to inspect and modify the system state during interaction with it, extending to all levels of access granted to the provider's own employees. It also governs the selection and involvement of independent experts and the procedural guarantees attaching to Article 101 penalties. Published on 21 July 2026, it entered into force on the twentieth day following publication, that is **10 August 2026** — after this cut-off.

This requires three questions to be kept apart, which version 1.1 treated as one:

1. **Does a European authority have the legal entitlement to obtain model access for evaluation?** From 2 August 2026, yes, and enforceably.
2. **Does it have the practical capacity to use it** — personnel, compute, published methods and an evaluator base? The reviewed record does not establish it. Key Action 1 dates European evaluation capacity to 2027, and the Plan records that most entities performing pre-deployment third-party evaluation are based outside the Union, which is a statement about where evaluators are concentrated rather than about their absence from Europe.
3. **Can European organisations more generally access advanced AI cyber capabilities?**
A separate question, addressed by the structured-access blueprint due in Q4 2026.

The distinction matters because collapsing the three would produce two symmetrical errors: treating a legal power as an operational capability, or continuing to describe evaluator access as constrained by provider discretion alone when it is now also a matter of enforceable obligation. Neither is accurate. The power was also untested at the cut-off, and its procedural arrangements, though published, were not yet in force.

Effect on the assessment. A statutory, sanction-backed route for Commission-led regulatory evaluation became applicable. That is narrower than general access to frontier models: it does not create access for operators, researchers, national authorities, independent evaluators or critical-sector organisations, and the structured-access blueprint addresses that separate question. No public evidence reviewed established a corresponding change in practical evaluation capacity. This change arrives on the AI Act's own statutory timetable and was not produced by the Plan.

³⁶ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence, OJ L, 2024/1689, 12 July 2024, ELI: <https://data.europa.eu/eli/reg/2024/1689/oj>, as amended by Regulation (EU) 2026/1744.[omnibus] Article 113 on entry into force and application; Article 91 on the power to request documentation and information; Article 92 on the power to conduct evaluations, in particular Article 92(3) on access through APIs or further appropriate technical means and tools including source code, Article 92(4) on the content of the request and Article 92(6) on implementing acts; Article 93 on the power to request measures; Article 101(1) on fines for providers of general-purpose AI models, in particular Article 101(1)(d).

³⁷ Commission Implementing Regulation (EU) 2026/1755 of 20 July 2026 on detailed arrangements for the conduct of certain proceedings by the Commission pursuant to Regulation (EU) 2024/1689 of the European Parliament and of the Council, OJ L, 2026/1755, 21 July 2026, CELEX 32026R1755, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1755>. Adopted on the basis of Articles 92(6) and 101(6) of Regulation (EU) 2024/1689. Recitals for the enumeration of access modalities, including application programming interfaces, internal access, source code, model weights, the infrastructure used for hosting the model, and access to inspect and modify the system state during interaction, extending to all levels of access granted to the provider's own employees; provisions on the selection and involvement of independent experts and on the procedural guarantees attaching to Article 101 penalties; Article 15 for entry into force on the twentieth day following publication, that is 10 August 2026.

6 What the Action Plan Sequences

6.1 The Plan's own schedule

The Plan is organised in three pillars — making frontier AI safe, accessible and deployable; applying cybersecurity fundamentals and accelerating patching; and building European capabilities — and it enumerates eight Key Actions, seven of which carry a target quarter.³⁸

Key Action	Content	Target
1	Commission to support establishment of an EU evaluation capacity for AI models, including cybersecurity	2027
2	Commission with ENISA to define a European Blueprint for structured access to advanced AI capabilities	Q4 2026
3	ENISA and the JRC to develop a secure testing platform for AI with advanced cyber capabilities	Q4 2026
4	ENISA to issue guidance, recommendations, advisories and best practices on AI-powered threats and secure AI integration	as of Q3 2026
5	Commission, Member States, ENISA and industry to make vulnerability-management practices and tools fit for the AI age	as of Q3 2026
6	ENISA to launch a first pilot of a Critical Open Source Resilience Campaign, accelerating patching including by leveraging AI	Q4 2026
7	Commission, with the European Cybersecurity Competence Centre and ENISA, to launch an EU Grand Challenge on AI-assisted vulnerability remediation	Q4 2026
8	Commission with Member States to aim to make AI Factories compute available to test, train and deploy models for cyber resilience on sovereign compute	none stated

Reporting this schedule changes what the readiness evidence means. A measure dated Q4 2026 with no service in August 2026 is not late; nor is an unexpired deadline evidence that delivery is on track, since no interim milestone is public against which progress could be measured. The absence of a stated deadline on Key Action 8, and its weaker verb — the

³⁸European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

Commission “will aim to make available” — is itself information about how firmly protected cyber compute is committed.

6.2 European 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.³⁹ The Plan’s addition is a commitment to strengthen European third-party assessment capacity, which the Commission expects to become operational in 2027.⁴⁰ The Plan notes that most entities currently performing pre-deployment third-party evaluation are based outside the Union.⁴¹

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

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

Readiness: G 2, R 1, T NE, A NE, U NE.

6.3 Structured access to advanced systems

The blueprint with ENISA responds to a stated problem: access to advanced cyber-capable AI “is increasingly governed by provider-specific and often non-European decisions”, and provider practice “often lacks transparency regarding the criteria applied” to decide which organisations obtain access.⁴² The blueprint is to be a guidance document specifying how providers can grant access to European organisations.

What the Plan leaves open is which users qualify, how deep the access goes, and whether providers will participate.

Additionality: new and Plan-specific.

Readiness: G 1, R NE, T NE, A NE, U NE.

6.4 ENISA–JRC secure testing platform

The platform would let organisations test AI cybersecurity applications in simulated environments, in cooperation with CERT-EU, Europol, sectoral authorities and Member States.⁴³

³⁹ European Commission, “EU AI Office Launches €9 Million Tender for Technical Support on GPAI Safety,” 10 July 2025, accessed 5 August 2026, <https://digital-strategy.ec.europa.eu/en/funding/eu-ai-office-launches-eu9-million-tender-technical-support-gpai-safety>.

⁴⁰ European Commission, “Governance and Enforcement of the AI Act,” accessed 5 August 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-act-governance-and-enforcement>.

⁴¹ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan’s pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

⁴² European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan’s pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

⁴³ European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan’s pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

The Plan specifies part of its operating model: participants use their own model access keys under shared security and confidentiality rules, each covering their own costs, using their own tools and retaining responsibility for their own use. Eligible-user criteria, technical depth and provider participation remain unspecified.

No accessible platform, service catalogue or completed test programme was identified. The target is Q4 2026.

Additionality: new and Plan-specific.

Readiness: G 2, R NE, T NE, A NE, U NE.

6.5 Defensive AI, vulnerability management and open source

Version 1.1 treated this cluster as redirected rather than Plan-specific, and found no evidence of a new common defensive service. Two pieces of evidence require that finding to be revised.

First, the Plan creates a named deliverable with a deadline: a first pilot of a Critical Open Source Resilience Campaign, led by ENISA with the Commission, Member States, open-source communities, Union entities and industry, to accelerate patching including by leveraging AI, targeted at Q4 2026 and intended to scale if successful.⁴⁴ Under this paper's own additionality rules that is **new and Plan-specific**, not redirected.

Second, guidance of the kind the Plan contemplates already existed at the Plan's launch. ENISA published *ENISA's view on Cybersecurity in the Frontier AI Era* on 7 July 2026: an initial set of operational recommendations for national competent authorities, policymakers, defenders and service providers, aimed at building the capabilities needed to face machine-speed threats.⁴⁵ Among its proposals are Union-wide benchmarks for the security evaluation of advanced models, including standardised testing against cyber ranges, exploitability metrics and simulations of chained attacks. ENISA states that the recommendations are not an all-inclusive checklist and will be refined, and it acknowledges input from the CSIRTs Network, EU-CyCLONe, its Advisory Group, its Cyber Partnership Programme, industry, the open-source community and academia.

The date matters for what may be concluded. The publication appeared on the same day as the Plan, so it **cannot** be evidence of the Plan's implementation. Its correct classification is a concurrent, Plan-aligned output: evidence that the type of guidance contemplated by Key Action 4 was already available when the Plan was adopted, not evidence that the Plan produced it. Version 1.1's finding that no ENISA output existed was wrong; the inference that this output demonstrates Plan delivery would be equally wrong.

Additionality: the open-source resilience pilot is new and Plan-specific; the wider vulnerability-management work is redirected; the guidance output is concurrent and Plan-aligned, with attribution to Plan implementation not independently established.

Readiness, open-source pilot: G 2, R NE, T NE, A NE, U NE.

Readiness, guidance output: G 2, R 1, T 1, A 2, U NE — the document is published and accessible to its intended readers, and no evidence of its use by them was identified.

⁴⁴European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

⁴⁵European Union Agency for Cybersecurity, *ENISA's View on Cybersecurity in the Frontier AI Era*, TLP:CLEAR, July 2026, https://www.enisa.europa.eu/sites/default/files/2026-07/ENISA%20view%20on%20cybersec%20in%20the%20frontier%20AI%20era_en_0.pdf, at the acknowledgements for CSIRTs Network and EU-CyCLONe input, and in the recommendations for Union-wide security-evaluation benchmarks including standardised cyber-range testing, exploitability metrics and chained-attack simulations. The document states the month of publication and not the day.

6.6 Grand Challenge

The Grand Challenge on AI-assisted vulnerability remediation is to bring together AI developers, cybersecurity companies, research organisations, critical-infrastructure operators and open-source communities, with support from the European Cybersecurity Competence Centre and in cooperation with ENISA, targeted at Q4 2026.⁴⁶ Rules, budget, outputs and the route from prototype to procurement or deployment are not public. An innovation competition may expand supply; it is not itself an operational defensive capability.

Additionality: new and Plan-specific.

Readiness: G 2, R NE, T NE, A NE, U NE.

6.7 AI Factories, Gigafactories and protected compute

AI Factories are inherited enabling infrastructure, politically reinforced by the Plan rather than created by it. The Plan positions existing AI Factories and future Gigafactories as part of a sovereign European infrastructure for AI and cybersecurity, and refers to a proposed Cloud and AI Development Act that would identify them as priority projects and allocate Union computing resources.⁴⁷

For general AI Factory services, Commission material supports network-level technical status and an access route, with the status divergence noted in section 4. No exercise, user output or repeated operational use was identified.

Readiness, general services: G 2, R 2, T 2, A C, U NE — access is marked conflicting because the two official descriptions differ.

For the protected, cyber-specific functions the Plan contemplates, the same infrastructure cannot be credited automatically. Key Action 8 makes the intention explicit but attaches no deadline and uses a weaker commitment verb. No public compute allocation or access rule for protected model or cyber evaluation was identified.

Readiness, protected cyber compute: G 1, R NE, T NE, A NE, U NE.

Gigafactories are intended to provide frontier-scale capacity. Official counts of the expression-of-interest response diverge — a dated update refers to 76 responses across 60 sites and 16 Member States, while the undated project page reports 77 proposals.^{48 49} The one-unit difference does not affect the assessment: the formal call and construction remained future steps.

Readiness, Gigafactories: G 2, R 1, T NE, A NE, U NE.

⁴⁶European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

⁴⁷European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

⁴⁸European Commission, "Commission Marks One Year of the AI Continent Action Plan with Two New Reports on AI Adoption and Policymaking," 9 April 2026, accessed 5 August 2026, <https://digital-strategy.ec.europa.eu/en/news/commission-marks-one-year-ai-continent-action-plan-two-new-reports-ai-adoption-and-policymaking>.

⁴⁹European Commission, "AI Gigafactories," accessed 5 August 2026, https://commission.europa.eu/topics/competitiveness/competitiveness-coordination-tool-projects/ai-gigafactories_en.

6.8 Reading the measures together

The pattern across the Plan's measures is not that they score badly. It is that evidence is concentrated in the mandate dimension and absent from the use dimension, which is what a set of commitments dated for later quarters should look like. The full five-dimension profile for all thirteen measures is in Appendix C, and it is a diagnostic rather than a measurement: coordination networks, response services, compute infrastructure, guidance documents and unbuilt platforms are not commensurable, and their totals are not compared.

What the profile does establish is the distribution of evidence stages. Inherited mechanisms carry reported implementation and, in two cases, evidence of a single exercise or deployment. Plan-specific measures carry mandates and, in one case, a procurement pathway that pre-dates the Plan. No measure carries evidence of repeated use.

7 Critical Infrastructure: Uneven Absorption

ENISA's NIS360 of May 2026 provides the comparative baseline. Its own scope statement sets the limit on what may be inferred: the assessment "covers the entire ecosystem of a sector", where a sector comprises its relevant actors including national authorities, rather than measuring individual organisations.⁵⁰ No statement about a particular operator follows from it.

NIS360 reports that maturity across sectors of high criticality has been steadily improving as organisations respond to policy requirements and to the threats they face, and that **banking, electricity and telecommunications remain the most mature and critical sectors**.⁵¹

Finance. Banking sits in the most mature group, and DORA contributes a structured regime for risk management and preparedness. Finance is therefore comparatively well positioned to integrate model evaluation, testing and AI-enabled threat information. The advantage is uneven by organisation size and is qualified by concentration in cloud, identity, payment, software and shared providers.

Electricity. Also in the most mature group. Its binding constraint is technical rather than institutional: long-lived operational-technology assets, strict availability and safety requirements, proprietary protocols and specialised engineering processes limit direct transfer of enterprise cyber tooling. The ENISA-JRC platform would have high value for this sector only if it provides realistic industrial scenarios and combined engineering and cyber expertise.

⁵⁰European Union Agency for Cybersecurity, *ENISA NIS360: Latest Insights in the Cybersecurity Maturity and Criticality of NIS Sectors of High Criticality*, May 2026, 82 pp., <https://doi.org/10.2824/8928447>, PDF at <https://www.enisa.europa.eu/sites/default/files/2026-05/ENISA%20NIS360%202026.pdf>. Page 6 for the scope statement that the assessment covers the entire ecosystem of a sector comprising its relevant actors including national authorities, for the finding that maturity has been steadily improving, and for banking, electricity and telecommunications remaining the most mature and critical sectors; the sector-by-sector overview at section 2, with energy at 2.1, finance at 2.4 and health at 2.5. The DOI returns HTTP 403 to automated clients because of publisher access controls; the report itself is served openly at the address above.

⁵¹European Union Agency for Cybersecurity, *ENISA NIS360: Latest Insights in the Cybersecurity Maturity and Criticality of NIS Sectors of High Criticality*, May 2026, 82 pp., <https://doi.org/10.2824/8928447>, PDF at <https://www.enisa.europa.eu/sites/default/files/2026-05/ENISA%20NIS360%202026.pdf>. Page 6 for the scope statement that the assessment covers the entire ecosystem of a sector comprising its relevant actors including national authorities, for the finding that maturity has been steadily improving, and for banking, electricity and telecommunications remaining the most mature and critical sectors; the sector-by-sector overview at section 2, with energy at 2.1, finance at 2.4 and health at 2.5. The DOI returns HTTP 403 to automated clients because of publisher access controls; the report itself is served openly at the address above.

Healthcare. Health is assessed separately in NIS360 and does not appear in the most mature group.⁵² High criticality combines 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.

Comparative summary. Version 1.1 published a five-dimension matrix of high, moderate and low ratings across governance, human capability, technical integration, procurement and remediation for each sector. That matrix is withdrawn. Fifteen cells were derived from a single ecosystem-level source, five of them as ranges, and none could be reproduced from the source at cell level. What the evidence supports is a three-way comparison, not a fifteen-cell measurement:

Sector	Position	Principal constraint	Basis
Finance	comparatively strong, uneven by organisation size	concentration in shared providers	ENISA maturity bands, ecosystem level
Electricity	institutionally strong, technically specialised	OT safety, legacy assets, integration	ENISA subsector assessment, ecosystem level
Healthcare	highly uneven and resource-constrained	fragmentation, legacy, staffing	ENISA maturity band and risk zone, ecosystem level

Cross-sector transmission. All three depend on telecommunications, cloud, data centres, managed ICT and security providers, identity services, software supply chains and electricity. A compromise in a common provider may propagate across sectors. Union-level capability must also pass through national authorities, sectoral supervisors, suppliers and operator systems, and failure at any interface can leave a formally available European mechanism with little effect on the ground.

8 Strategic Dependencies

Dependence is better assessed by function than by ownership. Five functions are informative: access, auditability, continuity, operational control and substitutability. A weak substitutability position does not mean no alternative exists; it means switching under operational stress may be slow, costly or technically disruptive.

Version 1.1 published a twenty-five-cell matrix rating five dependency categories against those five functions. **That matrix is withdrawn.** Fifteen of its cells had no cited source, two entire categories — advanced accelerators, and security software and open source — rested on no source in the reviewed record, and the paper itself described cell confidence as low to moderate. A matrix that cannot be reconstructed from its sources is a presentation of judgment as measurement.

What the evidence does support is narrower and, stated plainly, more useful.

⁵²European Union Agency for Cybersecurity, *ENISA NIS360: Latest Insights in the Cybersecurity Maturity and Criticality of NIS Sectors of High Criticality*, May 2026, 82 pp., <https://doi.org/10.2824/8928447>, PDF at <https://www.enisa.europa.eu/sites/default/files/2026-05/ENISA%20NIS360%202026.pdf>. Page 6 for the scope statement that the assessment covers the entire ecosystem of a sector comprising its relevant actors including national authorities, for the finding that maturity has been steadily improving, and for banking, electricity and telecommunications remaining the most mature and critical sectors; the sector-by-sector overview at section 2, with energy at 2.1, finance at 2.4 and health at 2.5. The DOI returns HTTP 403 to automated clients because of publisher access controls; the report itself is served openly at the address above.

Frontier models. The binding constraint is access, and it has two components that now diverge. For a Union regulator, access is a matter of enforceable obligation from 2 August 2026, under procedural arrangements published on 21 July and effective from 10 August, and untested at the cut-off. For European organisations generally, access remains governed by provider-specific and often non-European decisions whose criteria the Plan describes as lacking transparency.⁵³ Auditability and substitutability remain weak in both cases: models are not interchangeable, and evaluation depth depends on what a provider supplies.

Cloud infrastructure. The Commission's 2026 procurement awarded four contracts worth up to EUR 180 million and applied graduated sovereignty criteria; one consortium used Google Cloud-based technology through S3NS with services to be operated by European companies.⁵⁴ ⁵⁵ This supports a functional rather than binary reading of sovereignty. It does not demonstrate tested continuity or substitutability under stress, and one procurement cannot characterise European cloud dependence as a whole.

Advanced accelerators. No source in the reviewed record supports a rating. The dependence is widely asserted and, at this cut-off, not evidenced here. Recorded as an evidence gap rather than as a finding.

Skills and managed services. ENISA's *NIS Investments 2025* quantifies the constraint. Organisations across the Union report difficulty **attracting (76 per cent) and retaining (71 per cent)** cybersecurity professionals, intensified by a shortage of skilled candidates. Cybersecurity investment remains broadly at the level of the previous year — **9 per cent of IT budgets, median EUR 1.5 million** — but spending is increasingly directed towards technology and outsourcing rather than internal teams. Patching still takes months, and many organisations do not test their security.⁵⁶

That combination matters for how European capability reaches the ground. Shared services can compensate for a shortage of staff, but an organisation still needs expertise to request assistance, interpret results and remediate findings, and a budget increasingly spent on outsourcing does not build that expertise internally. The Cybersecurity Reserve shows that private capability can be operationalised through pre-arranged contracts and activation procedures; the same concentration creates common-mode risk.

The Plan's own technology-sovereignty framing links semiconductors, AI, cloud and open

⁵³European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

⁵⁴European Commission, "Commission Advances Cloud Sovereignty through Strategic Procurement," 17 April 2026, accessed 5 August 2026, https://commission.europa.eu/news-and-media/news/commission-advances-cloud-sovereignty-through-strategic-procurement-2026-04-17_en.

⁵⁵European Commission, "Sovereign Cloud Framework Explained," 1 June 2026, accessed 5 August 2026, https://commission.europa.eu/news-and-media/news/sovereign-cloud-framework-explained-2026-06-01_en.

⁵⁶European Union Agency for Cybersecurity, *NIS Investments 2025 — Main Report*, 38 pp., <https://doi.org/10.2824/7442427>, PDF at <https://www.enisa.europa.eu/sites/default/files/2026-02/NIS%20Investments%202025%20-%20Main%20report.pdf>. Page 5 for difficulty attracting (76 per cent) and retaining (71 per cent) cybersecurity professionals, for cybersecurity investment at 9 per cent of IT budgets with a median of EUR 1.5 million, and for spending increasingly focused on technology and outsourcing rather than internal teams; Insight 5 at page 21 for patching still taking months and for the proportion of organisations that do not test their security. The DOI returns HTTP 403 to automated clients; the report is served openly at the address above.

source, and treats AI Factories and Gigafactories as sovereign infrastructure.⁵⁷ ⁵⁸ On the evidence assembled here, Europe possesses substantial indigenous capacity while remaining constrained in access to frontier models and in specialist organisational capacity. The relevant objective is mission-relevant control and continuity, not autarky.

What this assessment does not establish. Advanced-accelerator dependence is widely asserted and is not evidenced in this paper: no source in the reviewed record supports a rating, and it is recorded as an evidence gap rather than as a finding. Cloud dependence is evidenced only through one procurement, which supports concern about continuity and substitutability without demonstrating failure under stress. Neither should be read as a verified concentration of European dependence.

9 Alternative Interpretations and Falsifiers

The central proposition is stated once, in section 1, and is not re-entered here as a competitor to itself. What follows is the case against it, in four readings.

9.1 Normal early-stage implementation

The strongest alternative. The observed pattern is mainly the predictable result of assessing a programme before its deadlines. The Plan was less than a month old at the cut-off, most deliverables were not yet due, and public institutions may hold personnel, compute access, methods or preparatory work that has not been disclosed. The procedural arrangements under Article 92(6) were published five days before the cut-off and took effect five days after it, so an applicable power was not yet an operating mechanism on 5 August, and was closer to becoming one than a reading of the Regulation alone would suggest. Existing AI Office procurement, ENISA and JRC expertise, EuroHPC infrastructure and private contractors could allow faster conversion than this assessment anticipates.

Under this reading the absence of public service evidence is normal programme timing, not evidence that complementary capacity was materially missing. This alternative is credible with moderate confidence, and it is the strongest of the four.

This is the reason the central proposition is confined to what the reviewed public record establishes, and the reason its confidence is moderate rather than high. What the paper reports is a transparency-bounded finding: the public record does not show the capacity. It does not show that the capacity is absent.

9.2 Regulatory layering

Prediction by 2028: guidance, calls and coordination accumulate without producing accessible services, faster remediation or measurable response improvement. *At the cut-off:* unresolved access criteria, uneven national implementation and many future deliverables. *Against:* the inherited layer is demonstrably usable rather than declaratory — the Reserve has an administering agency, funding, a procurement route, one deployment and one exercise — and the Plan's measures address gaps that are independently evidenced. *Falsifier:* repeated use of

⁵⁷European Commission, *Action Plan on Cybersecurity and Artificial Intelligence*, COM(2026) 577 final, Strasbourg, 7 July 2026, CELEX 52026DC0577, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52026DC0577>. Key Actions 1 to 8 are set out in the Plan's pillar sections, with a target quarter attached to seven of them and none to Key Action 8; the structured-access rationale, the testing-platform operating model, the open-source pilot, the Grand Challenge, the Cloud and AI Development Act reference and the sovereign-infrastructure framing are drawn from the same document.

⁵⁸European Commission, "Strengthening Europe's Tech Sovereignty," 3 June 2026, accessed 5 August 2026, https://commission.europa.eu/news-and-media/news/strengthening-europes-tech-sovereignty-2026-06-03_en.

Plan-specific services with measurable output. *Assessment*: a credible downside scenario, not demonstrated, and held with moderate confidence as a risk rather than as a description of the present.

9.3 Rapid operational convergence

Prediction: the combination of Article 92 powers, AI Office procurement, ENISA and JRC expertise and inherited cyber mechanisms allows the Union to operationalise evaluation and testing faster than this assessment expects. *At the cut-off*: complementary roles were assigned and cooperating entities named for the testing platform. *Against*: no procurement, platform, service catalogue, onboarding route or completed evaluation was identified for any deliverable dated Q4 2026 or 2027. *Falsifier for the negative*: a deliverable already accessible ahead of its quarter. *Assessment*: plausible but unsupported at the cut-off, and therefore entertained only with low confidence.

The ENISA guidance of 7 July 2026 is **not** evidence for this reading. It was published on the day the Plan was adopted, so it cannot demonstrate accelerated implementation of the Plan; it demonstrates that guidance of that type was already available.

9.4 Capability growth with continued external dependence

Prediction by 2028: European evaluation, compute and response capability expands while dependence on non-European models, cloud technology and specialist services persists. *At the cut-off*: expanding infrastructure and procurement coexist with constrained model access and untested substitutability, now partly offset in law by the change of 2 August 2026. *Against*: credible exit options, diversified supply and demonstrated continuity would weaken it. *Falsifier*: operational substitution during a material supplier disruption without loss of mission capability. *Assessment*: a plausible cross-cutting trajectory, compatible with either of the readings above. Moderate confidence applies to model access and specialist capacity, lower confidence to cloud substitution, and no assessment at all to accelerators, where this paper has no evidence.

9.5 Falsifiers of the historical assessment

The readings above concern the forward trajectory. A different class of evidence would revise the assessment of 5 August 2026 itself: evidence dated at or before the cut-off showing that a usable European evaluation capability already existed, meaning a completed third-party evaluation, allocated compute, published methods, an identified delivery structure and effective model access. None was identified within the search boundary stated in section 2.

10 Indicator Register

These are analytical verification thresholds, not official Union targets. They are anchored to the Plan's own quarters, so that a later assessment measures delivery against the schedule the Commission published rather than against an external expectation.

Indicator	Baseline at 5 August 2026	Verification threshold	Owner	Review
ENISA guidance (KA4)	first recommendations published July 2026	evidence of use: national implementation, operator adoption or follow-on advisories	ENISA	Jan 2027
Vulnerability management for the AI age (KA5)	no public output identified	published change to practices or tooling, with adoption route	Commission, ENISA, Member States	Jan 2027
Structured-access blueprint (KA2)	not published	adopted blueprint defining eligible users, access depth, safeguards and contingency arrangements	Commission, ENISA	Jan 2027
ENISA–JRC testing platform (KA3)	no accessible platform or catalogue	public service description, onboarding route, at least one enterprise and one OT-relevant environment	ENISA, JRC	Jan 2027
Open-source resilience pilot (KA6)	announced, no public rules	pilot launched with scope, participants and patching outcomes reported	ENISA	Jan 2027
Grand Challenge (KA7)	announced, no public rules	published rules, budget, outputs and pathway to procurement or deployment	Commission, ECCC	Jan 2027
European model-evaluation capacity (KA1)	no deployed service or completed output	published method, delivery structure and at least one completed third-party evaluation	AI Office	Jul 2027
Protected cyber compute (KA8)	no public allocation identified	documented compute allocation and access rules for protected model or cyber evaluation	EuroHPC, AI Office	Jul 2027
AI Act evaluation powers in use	powers applicable from 2 August 2026; procedural arrangements published 21 July 2026, effective 10 August 2026	at least one evaluation conducted under Article 92, with an identified evaluator and a published method	Commission, AI Office	Jan 2027
Cybersecurity Reserve	one deployment, one exercise	recurring reporting on activations, sectors, response time and service coverage	ENISA	Jul 2027
CRA reporting platform	reporting not yet mandatory	operational reporting with published volume or handling statistics	ENISA	Jan 2027

Indicator	Baseline at 5 August 2026	Verification threshold	Owner	Review
Sectoral absorption	no common participation and remediation dataset	sector-disaggregated evidence of onboarding, tests, exercises and remediation	ENISA, national authorities, operators	Jul 2027
Functional dependency management	no harmonised public register	documented exit, continuity or substitution tests for material suppliers	Commission, operators	Jul 2027

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

11 Conclusion

The Union's difficulty is no longer that its regulator lacks a legal route to the models it must assess. Since 2 August 2026 the Commission can compel access to a general-purpose AI model in order to evaluate it, refusal is a finable act, and since 21 July the detailed procedure for doing so has been published, taking effect on 10 August. The difficulty is that a power of this kind is only as useful as the apparatus behind it, and the reviewed public record does not show that apparatus: no completed evaluation, no published evaluation methodology, no identified delivery structure, no allocated compute for protected evaluation, and an evaluator base concentrated outside the Union. The Commission's own governance page places European evaluation capacity in 2027.

That is the conversion problem this assessment describes, and the Action Plan is best understood as a response to it. The Plan did not create the Union's cooperation, crisis-response, financial-resilience, product-security or compute architecture; those were in place before July 2026. What it adds is an order of work for the functions that were missing — evaluation capacity, structured access, secure testing, vulnerability remediation, open-source resilience, protected compute — distributed across quarters running to 2027, with the Indicator Register in section 10 stating what would count as delivery for each.

Two qualifications hold the conclusion in place. The first is temporal: four weeks is not long, most deliverables were not due, and an institution's silence is not its absence. The finding is that the public record does not establish the capacity, which is a weaker and more defensible claim than that the capacity is missing. The second is evidential: the negative half of the assessment rests on a targeted scan of official and institutional sources, not on an exhaustive search, and section 2 states that boundary.

The threat side gives the problem its urgency without settling it. Autonomous cyber capability is improving with inference-time compute and across model generations, but in bounded ranges that AISI describes as undefended and pre-breached. Nothing in that record demonstrates reliable operation against defended industrial systems, and treating benchmark progress as such a demonstration would repeat the error this paper is trying to avoid.

The assessment should be revised if evidence dated at or before the cut-off shows a usable European evaluation capability already in place: a completed third-party evaluation, allocated compute, a published method, an identified delivery structure and effective model access. It should be revised forward if the Plan's services become accessible and repeatedly used, with outputs linked to remediation or response, materially earlier than the quarters the Commission published.

Limitations

This assessment relies on public sources and cannot observe classified, internal or commercially confidential activity. The implementation update rests on the targeted official-source scan described in section 2.4, not on an exhaustive search of Member State, national authority or commercial records, and every negative finding is bounded by it. Absence of public evidence is not evidence that an internal capability does not exist, and this paper marks the difference explicitly rather than scoring both as zero.

Ten days separated the Plan's adoption from version 1.1's cut-off and four weeks separate it from this one, which prevents causal evaluation. The sectoral assessment used here evaluates ecosystems, not operators, so no statement about individual operators follows from it. Experimental benchmarks are run in undefended environments and may not transfer to live defended infrastructure; official benchmark pages revise individual figures, and this paper reports divergences rather than resolving them.

Readiness and sector positions are the author's assessments made against stated rules. They are not official ratings and do not substitute for organisation-level data.

The legal analysis in section 5 reads the Regulation and its implementing regulation directly, in the version in force at the cut-off, that is Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744. It is not legal advice, and the powers described were untested at the cut-off. The Digital Omnibus amends the third paragraph of Article 113 at points (a) and (c), which concern Chapters I and II and the high-risk timetable of Chapter III; it does not amend point (b), and it does not mention Chapter IX, Article 91, Article 92 or Article 101. The dating relied on here is therefore unaffected by it.

Appendix A. Measurement Rules

Dimensions apply only to operational mechanisms and service-like measures, never to legal regimes. Each is assigned from the notation in section 2.4 against these rules:

- institutional self-report may support mandate, resources or reported deployment; performance requires evidence of exercise, use or observable output;
- a future deadline or political commitment is not technical readiness;
- deployed infrastructure does not establish access to every intended service;
- an exercise or isolated deployment supports U 1, not repeated-outcome evidence;
- conflicting official descriptions are recorded as C, which lowers confidence rather than selecting a value;
- no dimension is scored optimistically where evidence is missing; NE is used.

Sensitivity. The comparison underlying Key Judgment 1 was tested against eleven alternative specifications, including removal of all self-reported operationality not independently corroborated, reduction of every uncertain positive by one level, exclusion of AI Factories, exclusion of the Cybersecurity Reserve, exclusion of both, and removal of the sector and dependency tables. The direction of the comparison holds in every specification in which it can be computed. It cannot be computed when all single-source claims are excluded, because every Plan-specific measure rests on the Action Plan alone.

That last result is a structural feature worth stating rather than a robustness finding: a communication is the primary source for its own contents, so a two-source rule cannot be satisfied for commitments it announces. Such claims are flagged in the claim ledger rather than corroborated artificially.

A further caution applies to the comparison itself. It places coordination networks, response services, compute infrastructure and unbuilt platforms on the same five dimensions, and inherited mechanisms have had years that Plan-specific measures have not. The comparison is robust and, taken alone, close to uninformative. The argument rests on the evidence-stage asymmetry and on the Plan's published quarters, not on any arithmetic gap.

Appendix B. Source Divergences

Three divergences between official sources were live at the research cut-off. Each is reported rather than resolved.

1. **AI Factory operational status.** A dedicated Commission page presents nineteen AI Factories and thirteen Antennas as operational and open to European users; a separate Commission overview states that work had begun on nineteen and that the majority were expected to be operational by the end of 2026. Both pages were re-read at this cut-off and both descriptions persist. This paper distinguishes network designation from site- and service-level operationality and records access as conflicting.^{59 60}
2. **Gigafactory expression-of-interest count.** A dated Commission update refers to 76 responses across 60 sites and 16 Member States; the undated project page reports 77 proposals. Because the project page carries no publication or update date, 77 is treated as an alternative current official figure rather than necessarily the later one.^{61 62}
3. **Human-effort estimate for the corporate cyber range.** AISI's April 2026 evaluation estimates that *The Last Ones* requires humans 20 hours to complete; the associated paper states roughly 14 hours for the same range. Both figures are recorded; neither is used as a precise measure of task difficulty.^{63 64}

Appendix C. Five-Dimension Readiness Profile

This profile is a supplementary diagnostic. It is not a measurement, the totals are not summed, and profiles for different object types are not compared. NA marks a dimension that does not

⁵⁹European Commission, "AI Factories," accessed 5 August 2026, <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>.

⁶⁰European Commission, "Artificial Intelligence," accessed 5 August 2026, https://commission.europa.eu/topics/artificial-intelligence_en.

⁶¹European Commission, "Commission Marks One Year of the AI Continent Action Plan with Two New Reports on AI Adoption and Policymaking," 9 April 2026, accessed 5 August 2026, <https://digital-strategy.ec.europa.eu/en/news/commission-marks-one-year-ai-continent-action-plan-two-new-reports-ai-adoption-and-policymaking>.

⁶²European Commission, "AI Gigafactories," accessed 5 August 2026, https://commission.europa.eu/topics/competitiveness/competitiveness-coordination-tool-projects/ai-gigafactories_en.

⁶³UK AI Security Institute, "Our Evaluation of Claude Mythos Preview's Cyber Capabilities," 13 April 2026, accessed 5 August 2026, <https://www.aisi.gov.uk/blog/our-evaluation-of-claude-mythos-previews-cyber-capabilities>, for the 32-step description of *The Last Ones*, the 20-hour human estimate, the first end-to-end completion in three of ten attempts, the average of 22 of 32 steps across attempts, and the next best model's average of 16 steps.

⁶⁴Linus Folkerts et al., "Measuring AI Agents' Progress on Multi-Step Cyber Attack Scenarios," arXiv:2603.11214v3, submitted 11 March 2026, current version 17 March 2026, <https://arxiv.org/abs/2603.11214>. Abstract and section 2 for the log-linear scaling of performance with inference-time compute with no observed plateau to 100 million tokens, the gain of up to 59 per cent between 10 and 100 million tokens, the generational rise from 1.7 to 9.8 average steps at a 10-million-token budget, the best single run of 22 of 32 steps against an estimated 14 human hours, and the industrial-range average of 1.2 to 1.4 steps of 7 with a maximum of 3; section 2 for the description of both ranges as sequential attack chains without active defenders or detection mechanisms.

apply to the object: a governance instrument has no technical readiness of its own, and a published document has no separate resource machinery beyond its own production.

Legal frameworks are excluded: DORA, NIS2, the AI Act and the CRA are baseline enablers whose dimensions are not commensurable with those of a service.

Measure	Object type	Additionality	G	R	T	A	U
CSIRTs Network	coordination network	inherited	2	2	2	2	1
EU-CyCLONe	crisis network	inherited	2	2	2	1	1
Cybersecurity Reserve	response service	inherited	2	2	2	2	1
CRA Single Reporting Platform	reporting service	inherited, developing	2	2	1	NE	NE
AI Factories, general services	compute/service network	inherited	2	2	2	C	NE
ENISA frontier-AI guidance	guidance output	concurrent, Plan-aligned	2	NA	NA	2	NE
Third-party model-evaluation capacity	evaluation service	new	2	1	NE	NE	NE
Structured-access blueprint	governance instrument	new	1	NA	NA	NE	NE
ENISA-JRC testing platform	testing service	new	2	NE	NE	NE	NE
Open-source resilience pilot	campaign	new	2	NE	NE	NE	NE
Grand Challenge	innovation instrument	new	2	NE	NE	NE	NE
Protected cyber compute	compute service	new	1	NE	NE	NE	NE
Gigafactories	infrastructure programme	inherited, reprioritised	2	1	NE	NE	NE

Notation. 2 established evidence; 1 partial, planned or pilot evidence; 0 affirmative evidence that the criterion is not met; NE no qualifying public evidence identified within the search boundary; NA not applicable to this object type; C conflicting official evidence.

On U. 2 requires repeated use by intended users with observable outputs or measured effects. 1 requires a single evidenced exercise, pilot or deployment. Availability of a measure to its users is scored under A and never under U: the publication of a document is A, and its use by readers is U. On this rule the guidance output scores A 2 and U NE.

Withdrawn from earlier versions. The R0 to R4 summary bands are not used. Under their own arithmetic a measure reached the top band, labelled operational evidence, on a single exercise; requiring repeated-outcome evidence for that band leaves no measure able to occupy it. A band that nothing can legitimately reach is not a measurement.

Bibliography

11.1 Union legislation and official documents

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

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 (DORA). OJ L 333, 27 December 2022. ELI: <https://data.europa.eu/eli/reg/2022/2554/oj>.

European Commission. Commission Implementing Regulation (EU) 2026/1755 of 20 July 2026 on detailed arrangements for the conduct of certain proceedings by the Commission pursuant to Regulation (EU) 2024/1689. OJ L, 2026/1755, 21 July 2026. CELEX 32026R1755. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1755>.

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) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI). OJ L, 2026/1744, 24 July 2026. CELEX 32026R1744. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1744>.

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>.

11.2 European Commission web sources

European Commission. “AI Factories.” Accessed 5 August 2026. <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>.

European Commission. “AI Gigafactories.” Accessed 5 August 2026. https://commission.europa.eu/topics/competitiveness/competitiveness-coordination-tool-projects/ai-gigafactories_en.

European Commission. “Artificial Intelligence.” Accessed 5 August 2026. https://commission.europa.eu/topics/artificial-intelligence_en.

European Commission. “Commission Advances Cloud Sovereignty through Strategic Procurement.” 17 April 2026. Accessed 5 August 2026. https://commission.europa.eu/news-and-media/news/commission-advances-cloud-sovereignty-through-strategic-procurement-2026-04-17_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 5 August 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. “EU AI Office Launches €9 Million Tender for Technical Support on GPAI Safety.” 10 July 2025. Accessed 5 August 2026. <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.” Accessed 5 August 2026. <https://digital-strategy.ec.europa.eu/en/policies/ai-act-governance-and-enforcement>.

European Commission. “Sovereign Cloud Framework Explained.” 1 June 2026. Accessed 5 August 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. Accessed 5 August 2026. https://commission.europa.eu/news-and-media/news/strengthening-europes-tech-sovereignty-2026-06-03_en.

11.3 European Union Agency for Cybersecurity

European Union Agency for Cybersecurity. *ENISA NIS360: Latest Insights in the Cybersecurity Maturity and Criticality of NIS Sectors of High Criticality*. May 2026. 82 pp. <https://doi.org/10.2824/8928447>. PDF: <https://www.enisa.europa.eu/sites/default/files/2026-05/ENISA%20NIS360%202026.pdf>.

European Union Agency for Cybersecurity. *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’s View on Cybersecurity in the Frontier AI Era*. TLP:CLEAR. 7 July 2026. https://www.enisa.europa.eu/sites/default/files/2026-07/ENISA%20view%20on%20cybersecurity%20in%20the%20frontier%20AI%20era_en_0.pdf.

European Union Agency for Cybersecurity. *NIS Investments 2025 — Main Report*. 38 pp. <https://doi.org/10.2824/7442427>. PDF: <https://www.enisa.europa.eu/sites/default/files/2026-02/NIS%20Investments%202025%20-%20Main%20report.pdf>.

European Union Agency for Cybersecurity. “Cyber Europe 2026: All Eyes on the EU’s Collective Response and Resilience.” Accessed 5 August 2026. <https://www.enisa.europa.eu/news/cyber-europe-2026-all-eyes-on-the-eus-collective-response-and-resilience>.

European Union Agency for Cybersecurity. “EU CyCLONe.” Accessed 5 August 2026. <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cyclone>.

European Union Agency for Cybersecurity. “EU Cybersecurity Reserve.” Accessed 5 August 2026. <https://www.enisa.europa.eu/topics/eu-incident-response-and-cyber-crisis-management/eu-cybersecurity-reserve>.

European Union Agency for Cybersecurity. “The EU Cybersecurity Reserve — FAQ.” Accessed 5 August 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. “Single Reporting Platform (SRP).” Accessed 5 August 2026. <https://www.enisa.europa.eu/topics/product-security-and-certification/single-reporting-platform-srp>.

11.4 Technical benchmarks and academic literature

Folkerts, Linus, et al. “Measuring AI Agents’ Progress on Multi-Step Cyber Attack Scenarios.” arXiv:2603.11214v3. Submitted 11 March 2026, current version 17 March 2026. <https://arxiv.org/abs/2603.11214>.

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, Onderai 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* (2026). <https://doi.org/10.1007/s11416-026-00622-3>.

UK AI Security Institute. “How Fast Is Autonomous AI Cyber Capability Advancing?” 13 May 2026. Accessed 5 August 2026. <https://www.aisi.gov.uk/blog/how-fast-is-autonomou-s-ai-cyber-capability-advancing>.

UK AI Security Institute. “Measuring AI Agents’ Progress on Multi-Step Cyber-Attack Scenarios.” Accessed 5 August 2026. <https://www.aisi.gov.uk/research/measuring-ai-agent-s-progress-on-multi-step-cyber-attack-scenarios>.

UK AI Security Institute. “Our Evaluation of Claude Mythos Preview’s Cyber Capabilities.” 13 April 2026. Accessed 5 August 2026. <https://www.aisi.gov.uk/blog/our-evaluation-of-cla-ude-mythos-previews-cyber-capabilities>.

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

Note on identifier access. The two ENISA reports carry Publications Office DOIs that return HTTP 403 to automated clients because of publisher access controls. Both DOIs resolve correctly through the Handle System, and both reports are served openly at the addresses above. The Union legal acts return HTTP 202 to automated clients through content negotiation on `data.europa.eu`; each was verified at the Publications Office CELLAR notice endpoint for official title, the act’s own date and the Official Journal reference. Neither condition is a citation defect.

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.

**How to check this edition**

Document AT-RR-2026-001, version 2.0, 2026-08-05.

DOI: 10.5281/zenodo.21810205.

Record: aletheia-technologies.it/verify/r/AVR-AT-RR-2026-001-v2.0-J3IOM/

Licence: CC BY 4.0.

Published by Aletheia Technologies.

Suggested citation

Paone, A. (2026). From Rules to Resilience: Assessing the EU's 2026 Action Plan on Cybersecurity and Artificial Intelligence. Version 2.0. DOI 10.5281/zenodo.21810205.

The record page states the exact bytes of this file, where it is deposited and which sources can be checked independently. It does not certify that the work is correct, original or first.