CYBERGUARD

AI-POWERED DEFENSE

WP6 · T6.4 POLICY DEVELOPMENT AND OPERATIONAL RECOMMENDATIONS

Strengthening SOC Capabilities through AI-Enabled Cyber Threat Intelligence and Operational Cooperation

Interim Policy Paper · Consultation Draft

Evidence-based recommendations for European cybersecurity actors



FOR CONSULTATION AT

CYBERGUARD European Cyber Resilience Forum

12 November 2026 · Regina Maria Ballroom, InterContinental Athénée Palace, Bucharest

Version 4.0 · November 2026 · Dissemination level PU – Public · GA 101190251

Co-funded by the European Union

DIRECTORATUL NAȚIONAL DE SECURITATE CIBERNETICĂ

ECCC EUROPEAN CYBERSECURITY COMPETENCE CENTRE

DOCUMENT CONTROL

Work package	WP6 — Dissemination, Communication and Exploitation
Task	T6.4 — Policy Development and Operational Recommendations
Task lead	DNSC, with I-ELINK and ENC
Contributes to	D6.2 — Exploitation activity, knowledge transfer and Policy recommendation (lead: BOLTON, due M36)
Dissemination level	PU — Public
Version	4.0 — Forum consultation edition
Date	November 2026 (project month 24 of 36)
Status	Consultation draft, presented at the CYBERGUARD European Cyber Resilience Forum, Bucharest, 12 November 2026

WHAT IS NEW IN VERSION 4.0

Version 3.1 set out the evidence and six measures. Version 4.0 turns the paper into the working document for the consultation required under Task 6.4. It now shows:

- which Forum sessions test which measure (Section 12.2);
- the 28 live-polling questions put to participants through **CYBERGUARD Quizy**, each linked to a measure (Section 12.3 and Annex B);
- structured slots where poll results and discussion outcomes are recorded after the Forum (Section 12.4);
- a traceability matrix extended with a Forum-input column (Annex A).

The evidence base and the six measures are unchanged from version 3.1. No consultation result is reported before the Forum has taken place.

STATUS NOTE

This document is a public consultation draft. The measures in Section 11 rest on findings recorded in the project's completed WP2 deliverables, summarised in Section 5 and traced in Annex A. Results from the WP4 adversarial defence work, the WP5 validation campaign and the Forum consultation will be incorporated into the final version delivered as part of D6.2.

Contents

At a glance 5
1. Executive Summary 6
2. Policy Context 7
2.1 The evidence base 7
2.2 The emerging EU institutional landscape for SOCs 7
2.3 CYBERGUARD's contribution 8
3. Legal and Policy Basis 9
4. Problem Statement 10
5. Evidence from CYBERGUARD 11
5.1 The threat landscape has shifted, and the shift is measurable 11
5.2 The defensive posture in the pilots is reactive, and the gaps are consistent 12
5.3 Threat intelligence is not arriving in a usable state 12
5.4 Willingness to share is constrained by compliance risk 12
5.5 Interoperability through layered normalisation 13
5.6 First pilot results 13
5.7 Evidence still outstanding 14
6. General Objective 14
7. Specific Objectives 14
8. Target Groups 15
9. Policy Options 15
Comparative assessment 16
10. Recommended Policy Option 16
11. Proposed Policy Measures 17
12. Consultation at the European Cyber Resilience Forum 19
12.1 Why the Forum 19
12.2 Programme and the measures each session tests 19
12.3 Live polling with CYBERGUARD Quizy 20
12.4 What we heard 21
12.5 How the input will be used 22
13. Implementation Approach 22
14. Monitoring Indicators 23
15. Expected Impact 24
16. Risks and Mitigation Measures 24
17. Conclusion 25
18. References 25
Annex A. Traceability Matrix 27

Annex B. Live-polling question set (CYBERGUARD Quizy)..... 28

At a glance



13	4	7	6
PARTNERS	COUNTRIES	PILOT ENVIRONMENTS	POLICY MEASURES

THE CYBERGUARD CONSORTIUM

Romania · 5 | Cyprus · 4 | Greece · 3 | Spain · 1 — partners per country

THE CENTRAL PROPOSITION

European cybersecurity actors should move **from reactive incident handling to proactive, risk-based and intelligence-led cyber defence**, using instruments already in force rather than waiting for new legal authority. The Cyber Solidarity Act has created the Cyber Hubs and the duty to interoperate; this paper proposes the operational content for that duty.

The six measures

M1 — CTI governance
Practical rules for collecting, prioritising, sanitising and sharing threat intelligence.

M2 — Interoperability
Structured formats and automated workflows from CTI to detection and response.

M3 — AI governance in the SOC
Human oversight, auditability and robustness for AI-supported analysis.

M4 — SME resilience
Simplified guidance and accessible CTI for smaller operators.

M5 — Trusted communities
Regular dialogue between hubs, SOCs, CSIRTs, authorities and industry.

M6 — Evidence traceability
Recommendations from EU projects traceable to the evidence behind them.

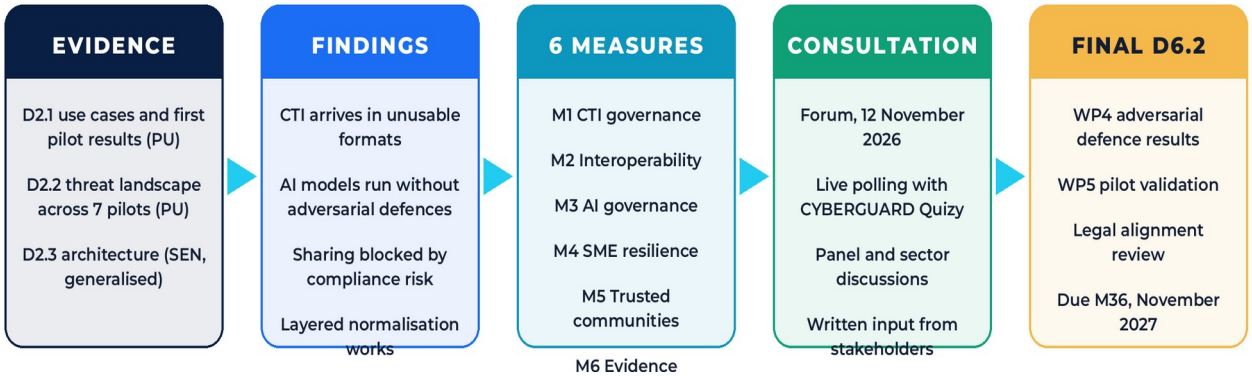


Figure 1 · From project evidence to the final recommendations in D6.2

1. Executive Summary

European cybersecurity operations face a threat environment that ENISA characterises as increasingly mature, fast-moving and difficult to attribute. In its 2025 Threat Landscape, ENISA analysed close to 4,900 confirmed incidents across the EU between July 2024 and June 2025: phishing accounted for roughly six in ten intrusions, vulnerability exploitation for approximately 21% of initial access, and newly disclosed flaws were weaponised within days of publication. Ransomware, though marginally down year-on-year, remained the most disruptive single category. [1]

Security Operations Centres, CSIRTs, public authorities, SMEs and critical infrastructure operators must absorb this volume and velocity under a regulatory framework that is itself consolidating rapidly. Three developments define the moment in which this paper is written:

- **The Cyber Solidarity Act** (Regulation (EU) 2025/38) entered into force on 4 February 2025, creating the European Cybersecurity Alert System and the first EU legal architecture for National and Cross-border Cyber Hubs: SOCs as instruments of Union policy rather than purely organisational assets. [2]
- **The AI Act simplification package** completed its passage in mid-2026, deferring high-risk obligations to 2 December 2027 (Annex III) and 2 August 2028 (Annex I) while leaving transparency duties and general-purpose AI enforcement live from 2 August 2026. [3][4]
- **The Action Plan on Cybersecurity and Artificial Intelligence** (COM (2026) 577 final, 7 July 2026) set out the Commission's coordinated response to advanced AI in cyber operations, framed as coordinating existing obligations under the AI Act, NIS2, DORA, the CRA and the Cyber Solidarity Act rather than creating a separate compliance regime. [5]

CYBERGUARD sits directly inside this architecture. The project is funded by the European Cybersecurity Competence Centre under the Digital Europe Programme, the same instrument and the same body responsible for building out the European Cybersecurity Alert System. Its technical work on AI-supported detection, malware analysis, CTI generation and sanitisation, and adversarial robustness addresses precisely the capability gaps that Cyber Hubs will need to close.

This paper translates that work into policy terms. It anchors each recommendation to a specific legal instrument or technical standard, so that the resulting policy direction is auditable, proportionate and directly usable by decisionmakers. The draft is put to practitioners, authorities and industry at the CYBERGUARD European Cyber Resilience Forum on 12 November 2026, where each measure is tested in a dedicated session and through live polling.

> **RECOMMENDED POLICY OPTION**
>
> **Option 3**: develop an integrated operational cooperation framework, implemented through six measures covering CTI governance, SOC interoperability, AI governance, SME resilience, trusted stakeholder communities, and evidence-based recommendation practice. Each measure is anchored to an instrument already binding on the actors concerned or entering into force during the project's lifetime.

2. Policy Context

European organisations are increasingly dependent on digital infrastructures, cloud services, interconnected systems and automated technologies. At the same time, cyber threats are becoming more complex, faster and harder to detect using traditional signature-based approaches.

2.1 The evidence base

The most recent ENISA Threat Landscape confirms three structural trends that shape this paper's priorities:

- **Volume and speed.** Vulnerability exploitation accounted for approximately 21% of initial-access intrusions, with adversaries weaponizing newly disclosed flaws within days of publication, well inside the window most organisations need to patch. [1]
- **AI on both sides of the contest.** AI-generated or AI-enhanced content featured in a substantial and rising share of observed phishing activity, industrialised through Phishing-as-a-Service platforms. The same automation defenders need is already available to attackers. [1][6]
- **Sectoral concentration.** Public administration, digital infrastructure, manufacturing and transport were among the most targeted sectors, closely matching the high-criticality sectors designated under NIS2 Annex I, which reinforces the case for aligning CTI and SOC cooperation with the Directive's scope. [1][7]

SOC and CSIRT teams are routinely required to process large volumes of alerts, indicators and threat intelligence feeds. Without adequate prioritisation, contextualisation and automation, this produces alert fatigue, delayed response and inefficient use of scarce analyst capacity.

2.2 The emerging EU institutional landscape for SOCs

Until recently, SOC cooperation in Europe was largely a matter of organisational initiative and bilateral trust. The Cyber Solidarity Act changes this. It establishes the European Cybersecurity Alert System as a network of National Cyber Hubs, designated by participating Member States and acting as the national reference point for collecting and analysing threat information, connected through Cross-border Cyber Hubs formed by three or more national hubs. [2]

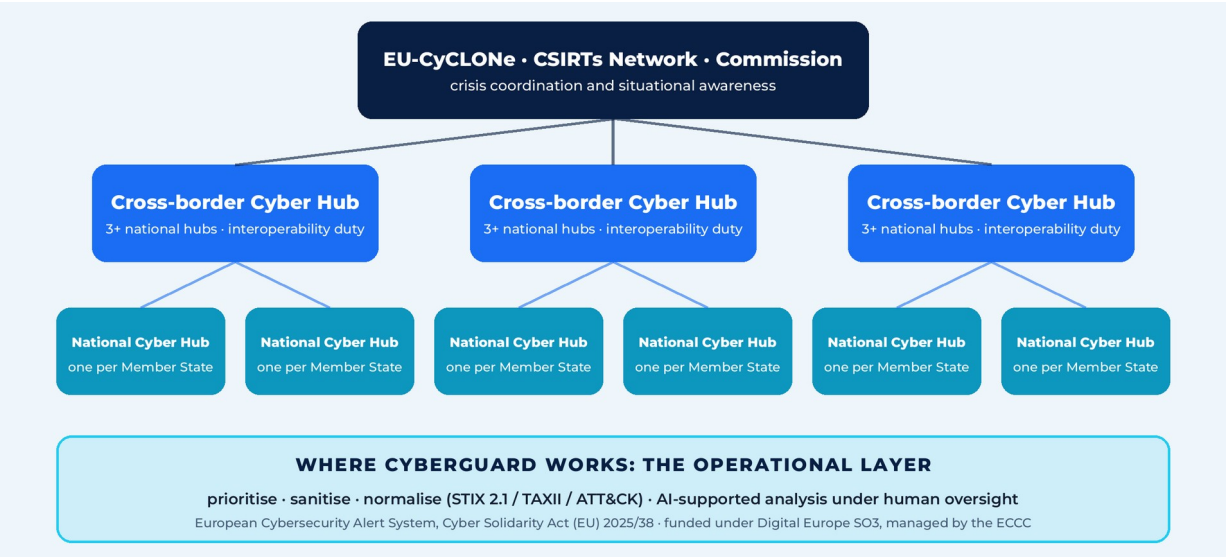


Figure 2 · The European Cybersecurity Alert System and the operational layer CYBERGUARD addresses

Three features of that architecture bear directly on the recommendations in this paper. First, Cross-border Cyber Hubs are required to ensure a high level of interoperability between themselves and to conclude cooperation agreements specifying information-sharing principles. Second, where they obtain information on a potential or ongoing large-scale incident, they must provide it to EU-CyCLONe, the CSIRTs Network and the Commission, in line with the crisis-management roles established by NIS2. Third, the hubs are expected to apply advanced technologies, artificial intelligence and data analytics prominently among them, to detect threats and share warnings across borders. [2]

In other words, the interoperability, trusted-sharing and AI-governance questions that CYBERGUARD addresses technically are now questions with a named institutional owner and a funding line. The Alert System is financed under Strategic Objective 3 (Cybersecurity) of the Digital Europe Programme and managed by the ECCC. [2] This paper is therefore addressed not only to Member State authorities in the abstract, but to a concrete and identifiable set of institutions currently being stood up.

2.3 CYBERGUARD's contribution

CYBERGUARD develops capabilities in:

- AI-supported threat detection and analysis
- malware analysis and extraction of indicators of compromise (IoCs)
- Cyber Threat Intelligence generation, sanitisation, normalisation and structured sharing
- SOC interoperability across organisational and national boundaries
- incident response support
- defence against adversarial attacks on machine-learning models and attacks targeting large language models
- pilot validation in critical and economically relevant sectors

This paper builds on those directions and translates technical results into practical recommendations for public authorities, SOCs, CSIRTs, SMEs and critical infrastructure operators, each cross-referenced to the EU legal instrument it must operate within.

This work is carried out by a consortium of thirteen partners established in four EU Member States, coordinated by I-ENERGYLINK, as shown in Figure 3.

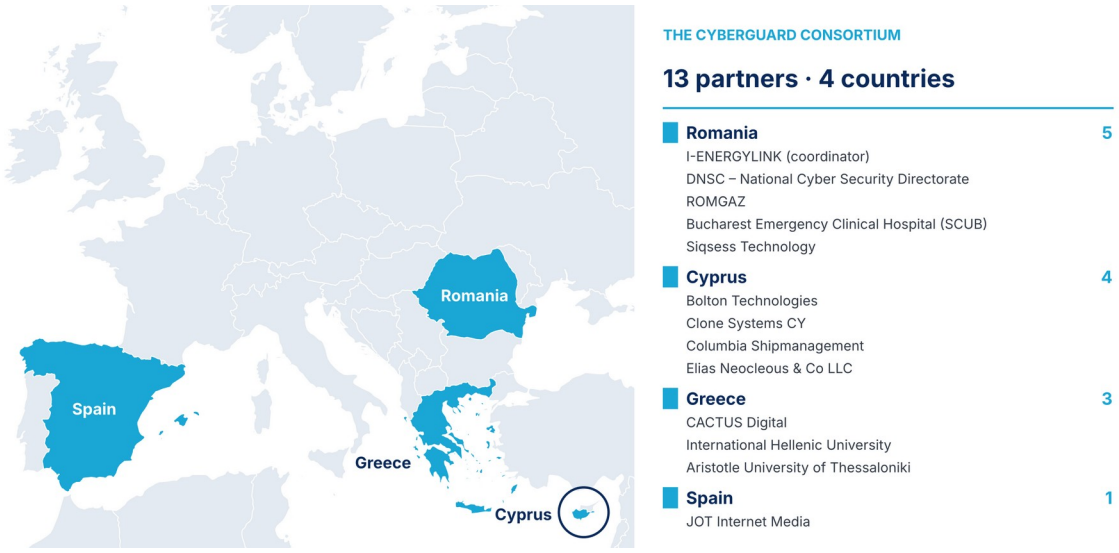


Figure 3 · The CYBERGUARD consortium: thirteen partners in Romania, Cyprus, Greece and Spain

3. Legal and Policy Basis

This paper is not formulated in a regulatory vacuum. Each recommended measure is designed to operate within, and reinforce, the existing EU cybersecurity acquis, so that adoption by Member States, SOCs and CSIRTs requires better operational use of instruments already in force rather than new legal authority.

Instrument	Type and status	Relevance to this paper
Cyber Solidarity Act (EU) 2025/38	Regulation; in force since 4 February 2025; implementation and ECCC calls ongoing	Establishes the European Cybersecurity Alert System, National and Cross-border Cyber Hubs, interoperability duties between hubs, and reporting to EU-CyCLONe and the CSIRTs Network. The primary institutional addressee of this paper.
NIS2 Directive (EU) 2022/2555	Directive; transposition deadline 17 October 2024; national implementation uneven	Art. 21 risk-management measures; Art. 23 incident-reporting cascade (24h / 72h / one month); Art. 29 voluntary information-sharing basis underpinning CTI governance; CSIRTs Network under Art. 15.
AI Act (EU) 2024/1689, as amended	Regulation; transparency and GPAI obligations applicable from 2 August 2026; high-risk obligations deferred to 2 December 2027 (Annex III) and 2 August 2028 (Annex I)	Articles 9–15 (risk management, data governance, logging, human oversight, cybersecurity resilience) frame the AI-governance measure for AI-supported SOC tooling.
Cybersecurity Act (EU) 2019/881	Regulation; in force	Mandates ENISA and establishes the EU cybersecurity certification framework, referenced for SOC tooling and CTI-platform assurance; amended by the Cyber Solidarity Act as regards managed security services.
Cyber Resilience Act (EU) 2024/2847	Regulation; phased application, main obligations from end-2027	Sets essential cybersecurity requirements for products with digital elements, relevant to SOC and CTI tooling procurement.

Instrument	Type and status	Relevance to this paper
DORA (EU) 2022/2554	Regulation; applicable since January 2025	Parallel ICT risk-management and incident-reporting regime for the financial sector; relevant where CYBERGUARD pilots touch financial-sector entities.
CER Directive (EU) 2022/2557	Directive; transposition ongoing	Complements NIS2 for the physical and operational resilience of critical entities engaged as target stakeholders.
GDPR (EU) 2016/679	Regulation; in force	Frames the data-protection safeguards required for CTI sharing, sanitisation and cross-border cooperation.
TLP — Traffic Light Protocol (FIRST.org)	International technical standard	Provides the classification scheme (RED / AMBER+STRICT / AMBER / GREEN / CLEAR) recommended for CTI governance and information sharing.
STIX 2.1 / TAXII 2.1 (OASIS); MISP	Open technical standards	Recommended structured formats for CTI interoperability between SOC, CSIRT and CTI-sharing platforms.

Policy context is provided by the Action Plan on Cybersecurity and Artificial Intelligence (COM (2026) 577 final), which pursues three objectives: promoting the safe use of advanced AI, strengthening EU cyber resilience, and expanding European AI capabilities for cybersecurity. It coordinates the instruments above rather than adding to them. [5] The recommendations in Section 11 are constructed to be consistent with that coordinating logic.

4. Problem Statement

Despite growing investment in cybersecurity tools and intelligence sources, many organisations still struggle to transform raw data into actionable intelligence. The recurring obstacles are fragmented security data, limited interoperability between tools, insufficient CTI governance, low trust in information sharing, a shortage of skilled analysts, and uncertainty about the operational use of AI in SOC environments under the AI Act's incoming obligations. [3][4]

SMEs and critical infrastructure operators are particularly exposed. ENISA's sectoral analysis shows manufacturing absorbing a disproportionate share of ransomware activity, and these organisations frequently lack the resources, in-house expertise or integrated tooling to process threat intelligence effectively or respond rapidly to incidents. [1][7]

CYBERGUARD's own assessment of its seven pilot environments confirms this picture and sharpens it: intelligence arrives in formats analysts cannot act on directly, AI models are deployed operationally without adversarial defences, and organisations withhold telemetry not through unwillingness but because they cannot assure compliance before data leaves their perimeter. Section 5 sets out these findings. [18]

The Cyber Solidarity Act sharpens rather than resolves this problem. It creates the obligation for Cyber Hubs to interoperate and to share but leaves open the operational question of how intelligence is prioritised, sanitised, quality-assured and rendered actionable, and the governance question of how AI-supported

analysis inside those hubs is kept accountable. That operational layer is where CYBERGUARD's technical work is positioned, and where this paper concentrates.

> **CORE PUBLIC POLICY PROBLEM**
>
> *How can European cybersecurity actors strengthen operational cyber resilience by improving the way threat intelligence, AI-supported analysis and SOC cooperation are governed, shared and used in practice, consistent with the Cyber Solidarity Act, NIS2, the AI Act and the Cybersecurity Act?*

5. Evidence from CYBERGUARD

The measures proposed in Section 11 rest on evidence generated by the project, not solely on the external threat-landscape literature summarised in Section 2. This section records the principal findings established to date.

Two sources carry most of the weight. D2.2, Comprehensive Analysis of the Threat Landscape: AI and Conventional Threats (delivered 28 February 2026, public), assessed the actual defensive posture of the seven pilot environments across the energy, healthcare, retail, maritime, digital marketing and cybersecurity sectors. D2.1, Detailed Use Cases Description and Scenarios for the Pilots (delivered 30 November 2025, public), documented the operational contexts and reported the first pilot results. D2.3, CYBERGUARD Architecture (delivered 6 June 2026), is classified SEN and is therefore drawn on here only at the level of generalised architectural lessons.

5.1 The threat landscape has shifted, and the shift is measurable

The project's threat analysis found a definitive move away from conventional IT and OT vulnerabilities towards an AI-enhanced landscape. Established vectors (multi-extortion ransomware, supply-chain compromise, credential-based initial access, exploitation of exposed services) remain prevalent across all pilots but are increasingly accelerated by artificial intelligence. At the same time, the introduction of AI inside the pilot infrastructures themselves opened new attack surfaces: adversarial machine-learning evasion, training-data poisoning, and prompt injection against large language models deployed in SOC workflows. [18]

One finding has direct consequences for CTI standardisation policy. When the project mapped its threat catalogue to the MITRE ATT&CK taxonomy, it found that several AI- and ML-specific threats, and several business-logic fraud patterns, are not directly captured by ATT&CK at all. The project recorded these gaps explicitly rather than forcing an approximate mapping. [18]

> **KEY FINDING**
>
> Policy that assumes existing taxonomies are sufficient to describe AI-era threats is assuming more coverage than currently exists.

5.2 The defensive posture in the pilots is reactive, and the gaps are consistent

The assessment of the seven pilot environments identified four systemic weaknesses recurring regardless of sector:

- **Over-reliance on signature-based detection.** Deployed SIEM and endpoint systems depend heavily on known indicators and static signatures and fail against activity with no historical profile. The retail pilot provided a concrete instance: AI-optimised memory-scraping malware that varies its execution timing specifically to avoid generating recognisable log patterns. [18]
- **AI and ML models deployed without adversarial defences.** Several pilot sectors had already adopted machine learning operationally, but without native protection against data poisoning, model evasion or prompt injection. In healthcare this was found to threaten diagnostic accuracy directly, if AI imaging tools were subtly manipulated. [18]
- **Alert fatigue and manual remediation.** Analysts face cognitive overload from alerts generated by rigidly tuned or outdated signatures, while the step from identifying a threat to mitigating it remains largely human. The resulting extension of mean time to respond gives machine-speed threats the window they need to move laterally, most severely in managed security environments where high alert volume meets limited staffing. [18]
- **Static, periodic security validation.** Offensive validation consists of scheduled manual penetration testing and vulnerability scanning, leaving blind spots between cycles and no continuous capability to simulate AI-enhanced attacker behaviour across converged IT/OT networks. [18]

5.3 Threat intelligence is not arriving in a usable state

This is the direct evidence base for Measure 1. Across the pilot infrastructures, CTI intake was found to rely heavily on non-machine-readable formats, text emails and standalone written reports, which blocks integration with SIEM tooling and forces analysts to extract tactical indicators by hand, introducing latency at precisely the point where speed matters most. [18]

Adoption of the standard taxonomies is inconsistent. Although STIX/TAXII and MITRE ATT&CK represent the industry norm for machine-readable intelligence, their use across the pilots was found to be uneven, and intelligence frequently arrived without contextual mapping to adversary tactics, techniques and procedures, limiting proactive threat hunting. The project also found that standard sandboxing regularly fails against environment-aware evasion, leaving a gap in the evasion-resistant analysis needed to extract usable indicators from novel malware. [18]

> **KEY FINDING**
>
> Mandating intelligence sharing does not by itself produce actionable intelligence. Format, structure and taxonomic mapping determine whether shared CTI reaches the analyst in a form they can act on.

5.4 Willingness to share is constrained by compliance risk

This is the evidence base for Measures 1, 2 and 5, and it qualifies the assumption behind information-sharing policy generally. Organisations in critical sectors, particularly healthcare (protecting patient data) and energy (protecting critical OT topologies), are deeply hesitant to share detailed incident telemetry because current

CTI sharing platforms lack the privacy-preserving capability, notably automated anonymisation and sanitisation, needed to guarantee compliance before data leaves the corporate perimeter. [18]

Where sanitisation cannot be assured, organisations default to withholding, and intelligence is shared retroactively, often long after a campaign has peaked. The project also found no centralised, trusted platform for cross-organisational visibility, producing operational silos, and identified a need for near-real-time, role-based sharing of anonymised attack patterns between National Coordination Centres and private-sector SOCs. Finally, manual intelligence-processing and reporting workflows were found insufficient to meet the incident-reporting timelines imposed by NIS2 and the Cyber Resilience Act. [18]

> **KEY FINDING**
>
> The barrier to sharing is technical and remediable, not cultural. Regulatory deadlines have outpaced the operational capability to meet them.

5.5 Interoperability through layered normalisation

This finding supports Measure 2 and is drawn from D2.3 at a level of generality appropriate to a public document. Interoperability across modules contributed by thirteen different partners rests on a small set of shared, standardised data representations rather than ad-hoc formats: common event formats at the edge, a normalised indicator and CTI model in the middle, and an ontology layer for knowledge-centric exploration, with threat intelligence represented in STIX 2.1, mapped to MITRE ATT&CK and exchanged over TAXII. This layering is what allows independently developed components to interoperate without bespoke point-to-point adapters, a directly transferable lesson for Cross-border Cyber Hubs facing the same problem at Member State scale. [19]

Two design decisions are relevant to Measure 3. Large-language-model serving is kept local by default, so that prompts and analysis content do not leave the deployment environment; and a guardrails gateway screens prompts and responses for personal data and secrets in both directions. Both show that AI-supported analysis and data-protection compliance are compatible when addressed architecturally rather than procedurally. [19]

5.6 First pilot results

Pilot activity began ahead of the contractual WP5 schedule, and the first results were reported in D2.1. In the digital marketing pilot, the anomaly-detection engine anticipated the advertising platform's own invalid-traffic adjustments with up to 85% overlap in early testing, identifying anomalous behaviour days before official metrics were updated, while AI-supported remediation guidance reduced manual review time by approximately 60%. [20]

up to 85%
overlap with the platform's own invalid-traffic adjustments, early testing

≈ 60%
less manual review time with AI-supported remediation guidance

These figures are early and indicative rather than final and should be presented as such. Their policy significance is not the percentages but the demonstration that detection methods developed for SOC environments transferred successfully into a non-traditional domain: evidence that investment in SOC capability yields returns beyond the SOC.

5.7 Evidence still outstanding

TO BE ADDED BEFORE THE FINAL VERSION IS ISSUED UNDER D6.2

D3.2 (Automated CTI Sanitizer and Normalizer, SEN): what the sanitisation work demonstrated about removing personal and attributable data from CTI at machine speed, and which residual judgements could not be automated. This directly tests the claim in Section 5.4 that the sharing barrier is remediable.

D3.1 (Advanced Malware Analysis Toolkit and CTI Feed System, SEN): findings on indicator quality, feed reliability and the practical cost of validation.

D1.3 (Ethics and Legal Guidelines, ENC-led): lessons from applying the Legal and Ethics Framework across a thirteen-partner, four-country consortium handling operational threat data.

D4.1 and D4.2 (LLM attack mitigation and adversarial defence): whether the defences hold under the adversarial conditions D2.2 identified.

D5.1, D5.2 and D5.3 (pilot deployment, evaluation and remediation guidance): consolidated validation results across all seven pilots.

Forum consultation, 12 November 2026: poll results and discussion outcomes (Section 12.4).

6. General Objective

To support stronger cyber resilience across Europe by improving the operational use of Cyber Threat Intelligence, AI-supported SOC capabilities and cross-sector cooperation, in alignment with the interoperability and information-sharing architecture established by the Cyber Solidarity Act, the risk-management and information-sharing obligations of NIS2, and the AI-governance requirements of the AI Act.

7. Specific Objectives

	Specific objective	Delivered by
SO1	Improve CTI actionability by encouraging prioritised, contextual and operationally relevant intelligence sharing, using TLP classification and STIX/TAXII structured formats.	M1
SO2	Strengthen SOC and CSIRT interoperability through common workflows, structured formats and better integration between CTI, detection and incident-response processes, consistent with the interoperability duties placed on Cross-border Cyber Hubs and the CSIRTs Network model under NIS2 Art. 15.	M2

	Specific objective	Delivered by
SO3	Support responsible use of AI in cybersecurity operations by promoting human oversight, transparency, accountability and compliance-by-design, aligned with AI Act Articles 9–15.	M3
SO4	Enhance cyber resilience of SMEs and critical sectors by facilitating access to practical threat intelligence, guidance and operational support, reflecting NIS2's extended scope to medium-sized entities.	M4
SO5	Promote trusted information sharing between public authorities, SOCs, CSIRTs, industry, research organisations and EU-funded cybersecurity projects, building on the NIS Cooperation Group, EU-CyCLONe and the emerging Cyber Hub network.	M5
SO6	Ensure that policy recommendations arising from EU-funded technical projects are traceable to operational evidence and validated with the communities expected to apply them.	M6

8. Target Groups

▸ National Cyber Hubs and Cross-border Cyber Hubs designated under the Cyber Solidarity Act	▸ National cybersecurity authorities and NIS2 competent authorities
▸ CSIRTs and SOC operators (national, governmental and sectoral)	▸ The European Cybersecurity Competence Centre and the National Coordination Centres
▸ Critical infrastructure operators under NIS2 and the CER Directive	▸ SMEs and business associations
▸ Public-sector institutions and cybersecurity service providers	▸ Research and academic organisations, and other EU-funded cybersecurity projects

9. Policy Options

OPTION 1

Maintain the current fragmented approach

Organisations continue to manage CTI, SOC workflows and incident response independently, using their own tools, processes and sharing practices, without a common alignment mechanism.

Advantages: Low implementation burden; no immediate structural change.

Disadvantages: Continued fragmentation, duplicated effort, limited trust, inconsistent intelligence quality, slower response, and weak alignment with both NIS2 Art. 29 and the interoperability duties the Cyber Solidarity Act now places on Cyber Hubs.

OPTION 2

Promote voluntary coordination and good practices

Organisations are encouraged to adopt common good practices for CTI sharing, AI-supported SOC workflows, data protection and stakeholder cooperation, without binding mechanisms.

Advantages: Flexible, easier to implement, adaptable to different maturity levels.

Disadvantages: Impact depends entirely on voluntary uptake and is likely to remain uneven across sectors and Member States, echoing the uneven NIS2 transposition experience.

OPTION 3 · RECOMMENDED

Develop an integrated operational cooperation framework

A structured framework for CTI sharing (using TLP and STIX/TAXII), SOC interoperability, AI governance aligned with the AI Act, and stakeholder engagement, designed to plug into the Cyber Hub architecture rather than sit alongside it, and supported by projects such as CYBERGUARD.

Advantages: Improves consistency, trust, interoperability and operational value; serves both technical and policy objectives; scales with organisational maturity; gives the Cyber Solidarity Act's interoperability requirements concrete operational content.

Disadvantages: Requires coordination, governance effort, stakeholder commitment, and continuous updating as AI Act and NIS2 implementing acts evolve.

Comparative assessment

Criterion	Option 1 · Status quo	Option 2 · Voluntary	Option 3 · Integrated
Regulatory alignment (Cyber Solidarity Act / NIS2 / AI Act)	Low	Medium	High
Implementation cost	Low	Low–Medium	Medium–High
Speed of impact	Not applicable	Slow, uneven	Gradual but consistent
Interoperability gain (STIX/TAXII, MISP)	None	Partial	Structural
SME accessibility	Low	Medium	Medium–High with tiered guidance
Resilience to fragmented adoption	Weak	Moderate	Strong, adaptable by maturity level

Green: strong · amber: partial · red: weak. Implementation cost is shown without colour, since a higher cost buys a larger effect.

10. Recommended Policy Option

The recommended option is Option 3: develop an integrated operational cooperation framework.

It offers the strongest added value because it connects technical innovation with practical operational needs and policy objectives, and because the institutional vehicle for it now exists. The Cyber Solidarity Act has created the hubs and imposed the duty to interoperate; what is missing is agreed operational content for that duty. Option 3 supplies it, while remaining anchored to instruments already binding on the actors concerned (the Cyber Solidarity Act, NIS2, the Cybersecurity Act) or entering into force during the project's lifetime (the AI Act, the Cyber Resilience Act).

The approach should remain flexible and adaptable, allowing organisations at different levels of maturity to adopt relevant elements progressively, consistent with NIS2's own differentiation between essential and important entities, and with the voluntary character of Member State participation in the Alert System.

11. Proposed Policy Measures

Each measure below states what it asks for, the instrument it rests on, the CYBERGUARD evidence behind it, and where it is tested at the Forum on 12 November 2026.

M1
CTI GOVERNANCE

Establish practical CTI governance principles

Develop guidance on how threat intelligence should be collected, prioritised, validated, sanitised, shared and used in operational environments, covering classification, confidence scoring, data protection and trust-based sharing arrangements.

ANCHOR TLP (FIRST.org) for classification; NIS2 Art. 29 as the legal basis for voluntary information-sharing arrangements; GDPR Arts. 6 and 9 safeguards for personal-data elements in CTI feeds; Cyber Solidarity Act information-sharing provisions for hub-to-hub exchange.

EVIDENCE D2.2 §4.2 (unstructured CTI intake; inconsistent STIX/TAXII adoption); D2.2 §4.3 (sanitisation as the precondition for sharing); D3.2 sanitiser results pending.

AT THE FORUM Panel 1 (11:05) · Panel 2 (12:30)

POLL QUESTIONS Q5, Q8, Q9, Q10, Q11, Q13, Q22, Q25 (Annex B)

M2
INTEROPERABILITY

Promote interoperability between SOC, CTI and incident-response workflows

Encourage structured formats, automated workflows and integrated platforms that allow CTI to move efficiently from analysis into detection and response, and give operational content to the interoperability obligations placed on Cross-border Cyber Hubs.

ANCHOR STIX 2.1 / TAXII 2.1 (OASIS) and MISP as recommended structured-exchange formats; Cyber Solidarity Act interoperability duties between Cross-border Cyber Hubs; the NIS2 CSIRTs Network cooperation model under Art. 15.

EVIDENCE D2.3 layered normalisation model (common event formats, normalised CTI schema, ontology layer); D5.2 pending.

AT THE FORUM ETSI keynote (10:15) · Panel 2 (12:30) · Technical snapshot (14:15)

POLL QUESTIONS Q10, Q12, Q22, Q23, Q25 (Annex B)

M3
AI GOVERNANCE IN THE SOC

Support AI governance in SOC environments

Define operational principles for responsible use of AI in cybersecurity, including human oversight, auditability, explainability, risk management, and safeguards against misuse of or over-reliance on automated decisions, recognising that Cyber Hubs are expected to deploy AI and data analytics at scale.

ANCHOR AI Act Arts. 9 (risk management), 10 (data governance), 12 (logging), 14 (human oversight) and 15 (cybersecurity resilience against data poisoning, adversarial inputs and model evasion); Action Plan on Cybersecurity and AI (COM(2026) 577 final).

EVIDENCE D2.2 §4.1 (ML models deployed without adversarial defences); D2.3 (local-by-default LLM serving; guardrails gateway screening for personal data and secrets); D4.1 and D4.2 pending.

AT THE FORUM Technical snapshot (14:15) · Policy session (15:45)

POLL QUESTIONS Q14, Q15, Q16, Q17, Q22, Q23, Q25 (Annex B)

M4
SME RESILIENCE

Strengthen SME-oriented cyber resilience

Develop simplified guidance, awareness materials and practical support mechanisms for SMEs, focusing on relevant threats, accessible CTI, incident preparedness and cooperation with authorities and service providers, without imposing compliance burdens disproportionate to SME resources.

ANCHOR NIS2's extended scope to medium-sized entities in scoped sectors; ENISA SME-oriented guidance under the Cybersecurity Act's awareness-raising mandate.

EVIDENCE D2.1 sector use cases; D2.2 sector-concentrated threat findings; pilot outcomes pending.

AT THE FORUM NCC-RO address (10:40) · Panel 1 (11:05) · Sector dialogue (14:45)

POLL QUESTIONS Q4, Q5, Q6, Q18, Q19, Q20, Q21, Q22, Q25 (Annex B)

M5
TRUSTED COMMUNITIES

Build trusted stakeholder communities

Support regular dialogue between Cyber Hubs, SOCs, CSIRTs, public authorities, critical infrastructure operators, SMEs, researchers and cybersecurity providers through workshops, exercises, policy events and structured consultations, plugging into existing fora rather than creating parallel structures.

ANCHOR The NIS Cooperation Group, the CSIRTs Network and EU-CyCLONe as existing EU-level cooperation fora; the Cyber Hub network as the emerging operational layer.

EVIDENCE D2.2 §4.3 (operational silos; retroactive sharing; need for NCC–SOC exchange); stakeholder engagement under WP6.

AT THE FORUM DNSC keynote (10:30) · Panel 2 (12:30) · Synergies (16:05)

POLL QUESTIONS Q9, Q11, Q12, Q13, Q21, Q22, Q25 (Annex B)

M6

EVIDENCE TRACEABILITY

Require traceability between operational evidence and policy recommendation

Establish the expectation that recommendations emerging from EU-funded cybersecurity projects are explicitly traceable to the operational evidence supporting them, and validated with the communities expected to apply them, so that project outputs can be assessed on their evidentiary basis rather than their plausibility.

ANCHOR European Commission Better Regulation Guidelines on impact assessment and stakeholder consultation.

EVIDENCE The traceability matrix at Annex A demonstrates the practice this measure recommends; the Forum consultation (Section 12) applies it.

AT THE FORUM Policy session (15:45) · Synergies (16:05)

POLL QUESTIONS Q7, Q22, Q24, Q25, Q26 (Annex B)

12. Consultation at the European Cyber Resilience Forum

12 NOVEMBER 2026 · REGINA MARIA BALLROOM, INTERCONTINENTAL ATHÉNÉE PALACE, BUCHAREST

12.1 Why the Forum

Task 6.4 commits the project to presenting the draft policy paper at a stakeholder event and gathering feedback before the final version is issued within D6.2. The CYBERGUARD European Cyber Resilience Forum is that event. It brings together public authorities, the National Coordination Centre, CSIRTs, SOC operators, critical infrastructure operators, SMEs, the cybersecurity industry, academia and other EU-funded projects, the same communities named as target groups in Section 8.

The programme is built so that every measure is examined by the people expected to apply it: regulators and legal experts on implementation (Panel 1), SOC and CTI practitioners on sharing and interoperability (Panel 2), project engineers on the technical evidence (technical snapshot), and operators from five sectors on what they actually need (sector dialogue).

12.2 Programme and the measures each session tests

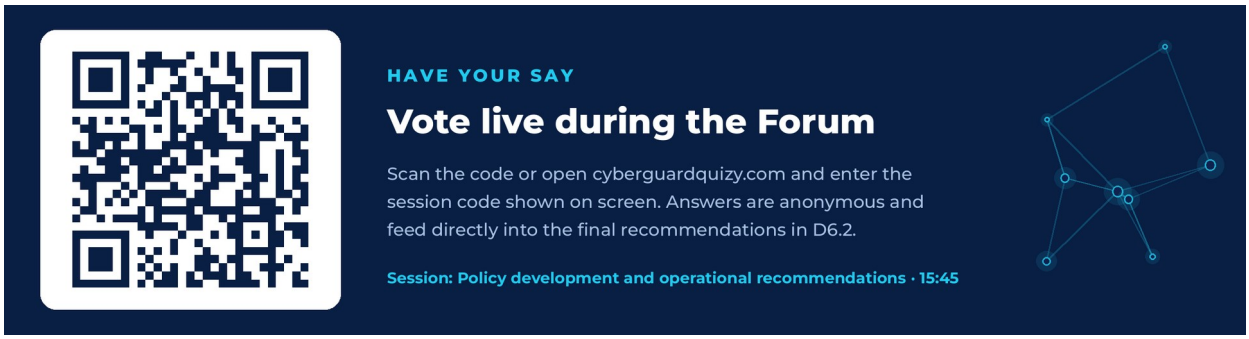
Time	Session	Measures tested
09:30	Opening remarks: CYBERGUARD and the European cyber resilience agenda	Framing
09:45	Special guest speaker (online)	Threat context
10:00	Ministerial address	Policy context
10:15	Institutional keynote: Strengthening Europe's SOC ecosystem through AI, CTI and cooperation (ETSI, online)	M2 · M3

Time	Session	Measures tested
10:30	Keynote: From European SOC innovation to collective cyber resilience (DNSC)	M2 · M5
10:40	Institutional address: From cybersecurity investment to operational resilience, the role of the NCC (NCC-RO)	M4 · M5
10:50	Strategic Keynote I: Cyber in a contested security environment	Threat context
11:05	Panel 1: From NIS2 and CRA to operational cybersecurity resilience	M1 · M4
12:05	Strategic Keynote II: Cyber resilience under fire, lessons from Ukraine	Threat context · M5
12:30	Panel 2: Cyber Threat Intelligence sharing and SOC interoperability	M1 · M2 · M5
14:05	CYBERGUARD overview: objectives, consortium, use cases and expected impact	Evidence
14:15	CYBERGUARD technical snapshot: malware analysis, CTI, AI-supported detection and incident response	M2 · M3
14:45	Sector dialogue: energy, healthcare, maritime, retail and digital services	M4
15:45	**Policy development and operational recommendations: first directions from CYBERGUARD · live polling**	All measures · M6
16:05	Synergies with other EU-funded cybersecurity projects	M5 · M6
16:45	Conclusions and next steps	Next steps

Times as in the final agenda of 5 November 2026.

12.3 Live polling with CYBERGUARD Quizy

During the policy session, participants answer 28 short questions on their phones through CYBERGUARD Quizy, the project's live-polling tool. Answers are anonymous. Each question is linked to a measure, so the results show where participants agree with the draft, where they would set different priorities, and what the draft has missed.



Theme	What it asks	Questions · measures
Who is in the room	Organisation type, sector, the single biggest obstacle	Q1–Q3

Theme	What it asks	Questions · measures
NIS2 and CRA in practice	Readiness, hardest requirements, what would help, regulatory coherence	Q4–Q7 · M4, M1, M6
Sharing threat intelligence	Frequency, barriers, formats, usefulness, who should lead, trust rules	Q8–Q13 · M1, M5, M2
AI in the SOC	Trust in AI triage, main concerns, level of human control, funding priorities	Q14–Q17 · M3
Sector resilience and SMEs	Most urgent sector, operator needs, SME support, best instrument	Q18–Q21 · M4, M5
Priorities for the final paper	Which recommendations to prioritise, sovereignty, uptake routes, one concrete measure, gaps	Q22–Q26 · M2, M3, M6
Closing	Europe's cyber resilience in 2030, usefulness of the day	Q27–Q28

The full question set with answer options is in Annex B.

12.4 What we heard

> **TO BE COMPLETED AFTER THE FORUM**
>
> *This section is deliberately left open. It will record the poll results and the main messages from each session once the Forum has taken place, and each finding will be linked to the measure it supports, qualifies or challenges.*

Poll results by theme

Theme	Responses	Main result and what it means for the measures
NIS2 and CRA in practice	[n]	[to be completed from the Quizy export]
Sharing threat intelligence	[n]	[to be completed from the Quizy export]
AI in the SOC	[n]	[to be completed from the Quizy export]
Sector resilience and SMEs	[n]	[to be completed from the Quizy export]
Priorities for the final paper	[n]	[to be completed from the Quizy export]
Closing	[n]	[to be completed from the Quizy export]

Key messages from the sessions

Session	Key messages	Effect on measure
Panel 1 · NIS2 and CRA to operational resilience	[to be completed from session notes]	M1 · M4: [supports / qualifies / new]

Session	Key messages	Effect on measure
Panel 2 · CTI sharing and SOC interoperability	[to be completed from session notes]	M1 · M2 · M5: [supports / qualifies / new]
Technical snapshot	[to be completed from session notes]	M2 · M3: [supports / qualifies / new]
Sector dialogue	[to be completed from session notes]	M4: [supports / qualifies / new]
Policy session and live polling	[to be completed from session notes]	All: [supports / qualifies / new]
Synergies with EU-funded projects	[to be completed from session notes]	M5 · M6: [supports / qualifies / new]

12.5 How the input will be used

- Poll results are exported from CYBERGUARD Quizy after the session and summarised by theme in Section 12.4.
- Session notes are taken by the WP6 team and checked with moderators before they are summarised.
- Every change to a measure that results from the consultation is recorded in Annex A, in the Forum-input column.
- Open-text answers (Q25 and Q26) are read in full; recurring proposals are reported, individual answers are not attributed.
- The consolidated version is shared with the consortium and the Executive Board before it is folded into D6.2.

13. Implementation Approach

The proposed approach should be implemented progressively:

- **Stage 1 · Evidence collection:** consolidate technical, operational, legal and stakeholder input from CYBERGUARD work packages, pilots, workshops and dissemination activities.
- **Stage 2 · Stakeholder consultation:** engage Cyber Hubs, SOCs, CSIRTs, SMEs, public authorities, industry and critical infrastructure operators to validate needs and expectations. The Forum of 12 November 2026 is the central consultation moment (Section 12).
- **Stage 3 · Drafting of recommendations:** transform collected evidence into policy and operational recommendations, distinguishing strategic recommendations for decision-makers from practical guidance for cybersecurity teams.

- **Stage 4 · Validation and refinement:** test recommendations with project partners and external stakeholders for clarity, feasibility and relevance, including a legal-compliance check against Cyber Solidarity Act, NIS2 and AI Act obligations.
- **Stage 5 · Dissemination and uptake:** publish and communicate the recommendations through CYBERGUARD channels, policy events and cooperation with the ECCC, the NIS Cooperation Group and the Cyber Hub network as appropriate.



Figure 4 · Road to the final recommendations. Dates after the Forum are planned.

14. Monitoring Indicators

The indicators below measure whether the recommended measures produce effect once adopted. They are intended for use by adopting organisations and by the bodies overseeing the Alert System, and are distinct from the project's own dissemination and engagement KPIs, which are tracked under D6.1.

Indicator	Type	Suggested target or method
Share of CTI received in structured format (STIX 2.1 / TAXII 2.1) by participating SOCs	**Output**	Baseline at adoption; target ≥ 70% within 24 months
Median time from CTI receipt to detection-rule deployment	**Outcome**	Measured per SOC; target: sustained reduction against own baseline
Proportion of shared CTI carrying explicit TLP designation and confidence scoring	**Output**	Target ≥ 90% of outbound sharing
Volume and reciprocity of cross-border CTI exchange between participating hubs	**Outcome**	Tracked bilaterally; reciprocity ratio as trust proxy
Documented human-oversight procedures for AI-supported triage and detection	**Output**	Present and current in all participating SOCs deploying AI tooling
Incidence of AI-tool decisions escalated to, and reversed by, human analysts	**Outcome**	Monitored as an over-reliance indicator; no fixed target
SME uptake of simplified CTI guidance and preparedness materials	**Outcome**	Survey-based; segmented by sector and entity size
Recommendations demonstrably traceable to operational evidence	**Output**	100% of recommendations issued under Measure 6
Perceived operational usefulness of shared CTI among receiving analysts	**Impact**	Periodic practitioner survey; complements volume metrics

15. Expected Impact

The proposed policy direction is expected to contribute to:

- improved operational use of Cyber Threat Intelligence, measurable through structured-format adoption and time-to-action
- stronger SOC and CSIRT cooperation, building on the CSIRTs Network and the emerging Cyber Hub architecture
- concrete operational content for the interoperability duties established by the Cyber Solidarity Act
- better preparedness against emerging threats, including AI-enabled phishing and rapid vulnerability exploitation
- reduced fragmentation in cybersecurity data and workflows
- stronger cyber resilience among SMEs and critical sectors, in line with NIS2's extended scope
- more responsible and transparent use of AI in SOC environments, consistent with AI Act Articles 9–15
- closer alignment between EU-funded technical innovation and European cybersecurity policy priorities, including the Action Plan on Cybersecurity and Artificial Intelligence

16. Risks and Mitigation Measures

Risk	Likelihood	Mitigation
Limited stakeholder engagement	**Medium**	Early stakeholder mapping and targeted outreach through existing networks: the NIS Cooperation Group, the CSIRTs Network and ECCC National Coordination Centres. The Forum consultation (Section 12) is the main engagement moment.
Overly technical recommendations	**Medium**	Separate technical guidance from policy messages; issue a decision-maker summary alongside the technical annex.
Low trust in information sharing	**Medium–High**	Promote clear data-handling and confidentiality rules using TLP classification and GDPR-compliant sanitisation, evidenced by the CYBERGUARD sanitisation work.
AI governance uncertainty as implementation dates shift	**Medium**	Frame human-oversight and compliance-by-design principles so they remain valid regardless of the final high-risk application dates.
Fragmented adoption across sectors and Member States	**Medium**	Adaptable recommendations for different maturity levels, mirroring NIS2's essential/important tiering and the voluntary character of Alert System participation.
Limited SME capacity	**High**	Simplified, practical, sector-relevant guidance avoiding burdens disproportionate to SME resources.
Regulatory drift (NIS2 implementing acts, Cyber Solidarity Act delegated acts, AI Act amendments)	**Medium**	Schedule a legal-alignment review before the final version is issued under D6.2, re-anchoring each measure to the instruments then in force.

Risk	Likelihood	Mitigation
Remaining WP4 and WP5 findings, or the Forum consultation, contradict or qualify the current evidence base	**Medium**	Annex A records the strength of evidence and the Forum input per measure, so any measure resting on incomplete or contested evidence is visible and can be revised rather than silently retained.

17. Conclusion

CYBERGUARD offers an opportunity to strengthen the connection between technical cybersecurity innovation and public policy at a moment when the European institutional architecture for SOC cooperation is being built rather than merely debated. The Cyber Solidarity Act has established the hubs and the duty to interoperate; the AI Act is settling the terms on which AI may be deployed inside them; the Action Plan on Cybersecurity and Artificial Intelligence has signalled the Commission's intention to coordinate rather than multiply these obligations.

This paper recommends an integrated operational cooperation framework supporting trusted CTI sharing, responsible use of AI, SOC interoperability and practical guidance for SMEs, CSIRTs, public authorities and critical infrastructure operators, designed to give operational content to obligations that already exist.

The final recommendations, to be issued as part of D6.2, will be evidence-based, stakeholder-validated and aligned with real operational needs, with each recommendation traceable to the CYBERGUARD evidence and to the Forum input supporting it. This consultation draft is put to the Forum to invite exactly that scrutiny.

18. References

[1] ENISA (2025). ENISA Threat Landscape 2025. European Union Agency for Cybersecurity. Analysis of confirmed incidents, 1 July 2024 – 30 June 2025.

[2] Regulation (EU) 2025/38 of the European Parliament and of the Council 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). In force 4 February 2025.

[3] Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, 12 July 2024.

[4] Digital Omnibus on AI: targeted amendments to Regulation (EU) 2024/1689. Political agreement 7 May 2026; European Parliament endorsement 16 June 2026; Council adoption 29 June 2026. Deferral of high-risk obligations to 2 December 2027 (Annex III) and 2 August 2028 (Annex I).

[5] European Commission (2026). Action Plan on Cybersecurity and Artificial Intelligence, COM(2026) 577 final, 7 July 2026. See also press release IP/26/1544.

[6] ENISA (2025). Sectorial Threat Landscape reports, including Public Administration.

[7] Directive (EU) 2022/2555 of the European Parliament and of the Council on measures for a high common level of cybersecurity across the Union (NIS2), OJ L 333, 27 December 2022. Transposition deadline 17 October 2024.

[8] Regulation (EU) 2019/881 of the European Parliament and of the Council on ENISA and on information and communications technology cybersecurity certification (Cybersecurity Act), OJ L 151, 7 June 2019.

[9] Regulation (EU) 2024/2847 of the European Parliament and of the Council on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act), OJ L, 20 November 2024.

[10] Regulation (EU) 2022/2554 of the European Parliament and of the Council on digital operational resilience for the financial sector (DORA), OJ L 333, 27 December 2022.

[11] Directive (EU) 2022/2557 of the European Parliament and of the Council on the resilience of critical entities (CER), OJ L 333, 27 December 2022.

[12] Regulation (EU) 2016/679 of the European Parliament and of the Council on the protection of natural persons with regard to the processing of personal data (GDPR), OJ L 119, 4 May 2016.

[13] FIRST.org. Traffic Light Protocol (TLP): Standard Definitions and Usage Guidance, version 2.0.

[14] OASIS Open. STIX 2.1 and TAXII 2.1 Specifications: structured threat-information exchange standards.

[15] MISP Project. MISP: Open Source Threat Intelligence Platform documentation.

[16] European Cybersecurity Competence Centre. Calls for Expressions of Interest: National and Cross-border Cyber Hubs, 2025 and 2026 rounds.

[17] European Commission. Better Regulation Guidelines: impact assessment and monitoring methodology.

[18] CYBERGUARD (2026). D2.2: Comprehensive Analysis of the Threat Landscape: AI and Conventional Threats, v1.0. Dissemination level: PU – Public. Delivered 28 February 2026.

[19] CYBERGUARD (2026). D2.3: CYBERGUARD Architecture, v1.0. Dissemination level: SEN – Sensitive. Delivered 6 June 2026. Cited here only at the level of generalised architectural findings.

[20] CYBERGUARD (2025). D2.1: Detailed Use Cases Description and Scenarios for the Pilots, v1.0. Dissemination level: PU – Public. Delivered 30 November 2025.

NOTE ON SOURCING

All statistical claims in this paper are drawn from primary ENISA publications and all legal claims from the Official Journal texts. Figures cited from the ENISA Threat Landscape should be verified against the published report before the final version is issued, and the reference list updated with the exact publication date and page references.

Annex A. Traceability Matrix

This annex demonstrates the practice recommended in Measure 6: each proposed measure is linked to the CYBERGUARD deliverable providing its evidence base, the strength of that evidence is stated rather than implied, and the input received at the Forum is recorded alongside. Entries recorded as outstanding will be completed before the final version is issued under D6.2.

Measure	Evidence source	Deliverable status	Evidence status	Forum input (12 Nov 2026)
M1 · CTI governance	D2.2 §4.2 unstructured intake, inconsistent STIX/TAXII adoption; D3.2 sanitiser	D2.2 delivered 28 Feb 2026 (PU); D3.2 SEN, pending	Substantiated; sanitiser results outstanding	*[to be completed]*
M2 · Interoperability	D2.3 layered normalisation model; D5.2 SOC resilience toolkit	D2.3 delivered 6 Jun 2026 (SEN); D5.2 pending	Substantiated at architectural level	*[to be completed]*
M3 · AI governance	D2.2 §4.1 undefended ML models; D2.3 local LLM serving and guardrails gateway; D4.1, D4.2	D2.2 and D2.3 delivered; WP4 pending	Substantiated; adversarial validation outstanding	*[to be completed]*
M4 · SME resilience	D2.1 sector use cases; D2.2 sector-concentrated threat findings	Both delivered (PU)	Substantiated; pilot outcomes outstanding	*[to be completed]*
M5 · Trusted communities	D2.2 §4.3 absence of trusted shared platform and retroactive sharing	Delivered (PU)	Substantiated	*[to be completed]*
M6 · Evidence traceability	This annex; Forum consultation (Section 12)	Current	Complete for v4.0	*[to be completed]*

Prepared under Task 6.4 by DNSC with I-ELINK and ENC, for inclusion in Deliverable D6.2 (lead: BOLTON).

CYBERGUARD has received funding from the European Union under the Digital Europe Programme, managed by the European Cybersecurity Competence Centre, under Grant Agreement No. 101190251. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the granting authority.

Annex B. Live-polling question set (CYBERGUARD Quizy)

The 28 questions put to participants during the policy session, grouped by theme, with question type, answer options and the measure each one informs. Rating questions use a 1–5 scale; word-cloud and open questions collect free text.

WHO IS IN THE ROOM

#	Question and options	Type	Informs
Q1	**Which best describes your organisation?** National authority / regulator · CSIRT / SOC (public) · Essential or important entity (NIS2) · Cybersecurity vendor / MSSP · SME (non-security) · Research / academia · EU-funded project / EU body · Other	Single choice	Profile
Q2	**Which sector do you mainly work in?** Energy · Healthcare · Maritime / transport · Retail · Digital services / ICT · Public administration · Finance / insurance · Other	Single choice	Profile
Q3	**In one word: what is the biggest obstacle to cyber resilience in your organisation today?**	Word cloud	Framing

NIS2 AND CRA IN PRACTICE

#	Question and options	Type	Informs
Q4	**How ready is your organisation for its NIS2 obligations? (1 = not at all, 5 = fully ready)**	Rating 1–5	M4
Q5	**Which NIS2 / CRA requirements are hardest to put into practice? (choose up to 3)** Risk management measures (Art. 21) · Incident reporting within 24h / 72h · Supply-chain and supplier security · Vulnerability handling and disclosure · CRA product security requirements and conformity · Management body accountability and training · Overlapping obligations from several regulations	Multiple choice	M1 · M4
Q6	**What would help you most to meet NIS2 and CRA obligations?** Clear national guidance and templates · One single reporting point for all regulations · Sector-specific implementation guides · Funding for tools and SOC services · Shared services for SMEs (e.g. SOC-as-a-service) · Training and skills	Single choice	M4
Q7	**Do NIS2, the CRA and other EU rules feel coherent and complementary in your daily work? (1 = very fragmented, 5 = fully coherent)**	Rating 1–5	M6

SHARING THREAT INTELLIGENCE

#	Question and options	Type	Informs
Q8	**How often does your organisation share cyber threat intelligence (CTI) with others?** Never · Only when an incident happens · Occasionally (monthly) · Regularly (weekly or more) · Automatically, in real time	Single choice	M1
Q9	**What stops you from sharing more CTI? (choose up to 3)** Legal uncertainty (GDPR, confidentiality, liability) · Lack of trust in who receives the data · No machine-readable formats / tools (STIX/TAXII, MISP) · Lack of staff and time · Low quality or relevance of what we receive · No clear benefit for our organisation · Internal approval procedures	Multiple choice	M1 · M5

#	Question and options	Type	Informs
Q10	**Which formats do you use to receive or share CTI? (choose all that apply)** Machine-readable (STIX/TAXII, MISP) · Structured but proprietary (vendor feeds) · E-mail, PDF or reports · Phone / personal contacts · We do not use CTI	Multiple choice	M1 · M2
Q11	**How useful is the threat intelligence you currently receive from national or EU sources? (1 = not useful, 5 = very useful)**	Rating 1–5	M1 · M5
Q12	**Who should play the central role in CTI sharing in Romania?** DNSC / national CSIRT · Sector CSIRTs and ISACs · National Coordination Centre (NCC-RO) · Cross-border European SOCs (Cyber Solidarity Act) · Private MSSPs and vendors · A mix, with clear roles for each	Single choice	M2 · M5
Q13	**Which rules would make you trust a CTI sharing platform? (choose up to 3)** Clear legal basis and data-protection rules · Traffic Light Protocol (TLP) respected and enforced · Anonymisation of the source · Independent governance and audit · Reciprocity: you get back what you give · EU-hosted, sovereign infrastructure	Multiple choice	M1 · M5

AI IN THE SOC

#	Question and options	Type	Informs
Q14	**How much do you trust AI-supported tools to triage alerts and support incident response in a SOC? (1 = not at all, 5 = fully)**	Rating 1–5	M3
Q15	**What are your main concerns about AI in security operations? (choose up to 3)** Wrong or invented results (hallucinations) · Attacks against the AI itself (poisoning, prompt injection) · Lack of explainability · Data protection and confidentiality · Dependence on non-EU providers · Loss of analyst skills · Unclear liability and AI Act obligations	Multiple choice	M3
Q16	**What level of human control should AI have in a SOC?** AI only suggests, analysts decide everything · AI acts on low-risk tasks, analysts approve the rest · AI acts autonomously, analysts review afterwards · Full automation where possible	Single choice	M3
Q17	**Which AI-supported capabilities should EU and national funding prioritise? (choose up to 3)** Malware analysis and classification · Alert triage and prioritisation · CTI correlation and enrichment · Anomaly detection in OT / industrial systems · Automated detection rules (YARA, Sigma) · Report summarisation for decision makers · Testing AI robustness (red teaming)	Multiple choice	M3

SECTOR RESILIENCE AND SMES

#	Question and options	Type	Informs
Q18	**Which sector needs the most urgent support for cyber resilience? (one or two words)**	Word cloud	M4
Q19	**What do operators in your sector need most? (choose up to 3)** 24/7 monitoring they can afford · OT / ICS security expertise · Sector-specific threat intelligence · Incident response support · Tested exercises and playbooks · Secure-by-design products from suppliers · Skilled people	Multiple choice	M4

#	Question and options	Type	Informs
Q20	**How well are small and medium-sized operators supported today to meet cyber requirements? (1 = not at all, 5 = very well)**	Rating 1–5	M4
Q21	**Which instrument would most improve operational resilience in your sector?** Shared or sector SOC services · Pilots and testing environments (like CYBERGUARD) · Vouchers / grants for SMEs · Mandatory minimum security baseline · Public-private information sharing groups (ISACs) · Joint exercises and simulations	Single choice	M4 · M5

PRIORITIES FOR THE FINAL PAPER

#	Question and options	Type	Informs
Q22	**Which policy recommendations should CYBERGUARD prioritise? (choose up to 3)** Machine-readable, interoperable CTI by default · A legal safe harbour for good-faith CTI sharing · Security and robustness requirements for AI used in SOCs · One reporting point across NIS2, CRA and GDPR · Shared SOC services for SMEs and smaller operators · Sector-specific guidance and pilots · Skills and training programmes for SOC analysts · Using EU funds to scale tested solutions	Multiple choice	All
Q23	**How important is EU technological sovereignty (European tools, EU-hosted data) for SOCs? (1 = not important, 5 = critical)**	Rating 1–5	M2 · M3
Q24	**How can EU-funded projects like CYBERGUARD best turn results into real use?** Open-source tools and reusable components · Free pilots with national and sector operators · Integration into national and European SOC networks · Training and documentation for operators · Links with procurement and follow-up funding	Single choice	M6
Q25	**What is ONE concrete measure you would ask policymakers to adopt in the next 12 months?**	Open text	All
Q26	**Is there anything missing from today's discussion that the CYBERGUARD policy recommendations should address?**	Open text	M6

CLOSING

#	Question and options	Type	Informs
Q27	**One word to describe Europe's cyber resilience in 2030**	Word cloud	Framing
Q28	**How useful was today's Forum for your work? (1 = not useful, 5 = very useful)**	Rating 1–5	Evaluation