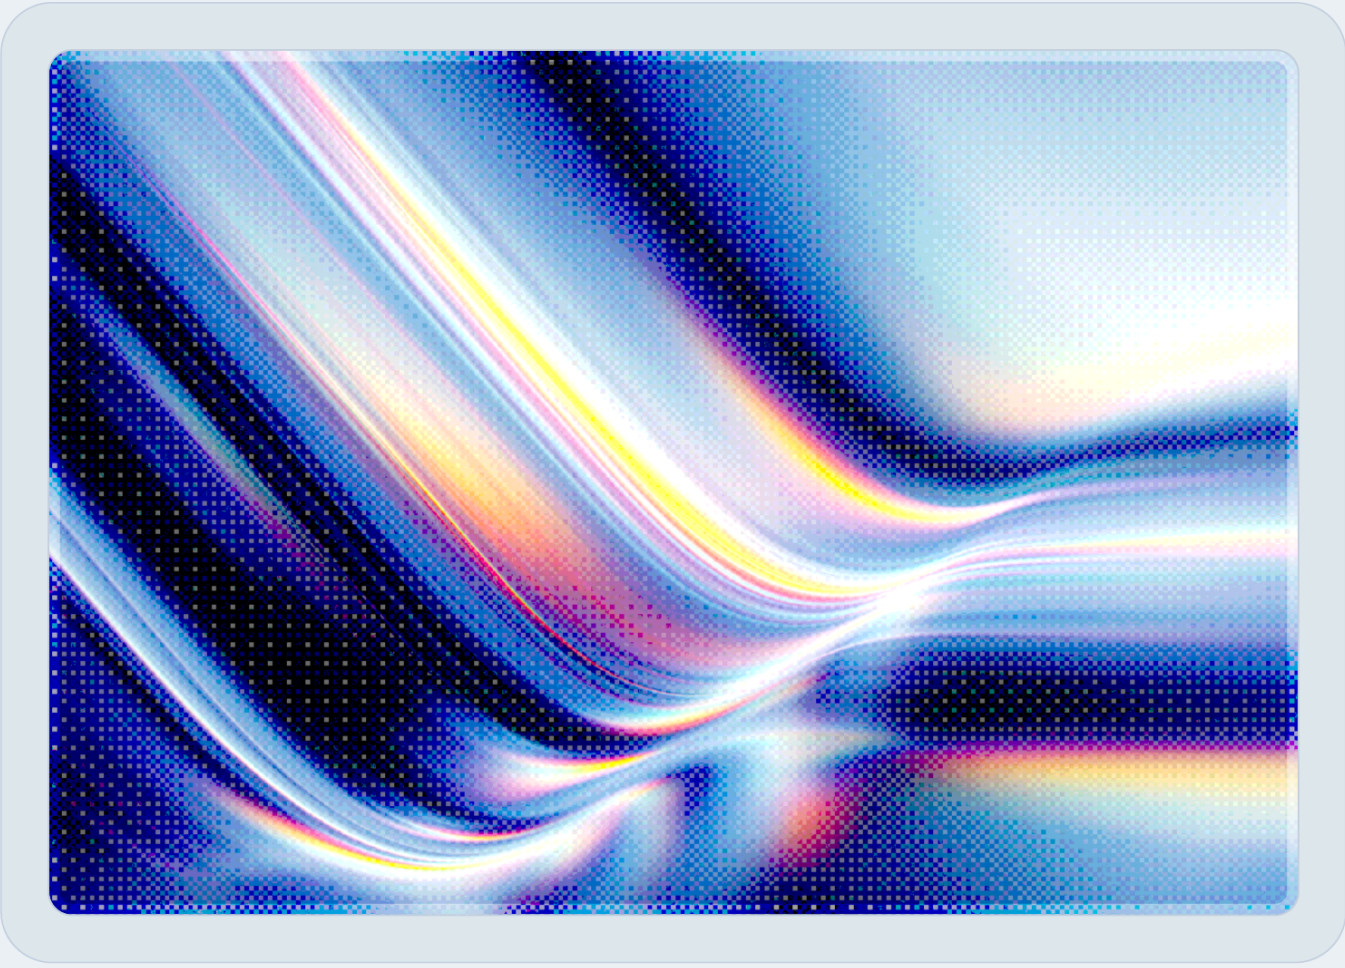
CSAI Foundation | Cloud Security Alliance

AI-Industrialised Vulnerability Discovery: ENISA 2026 and NIS2/CRA Compliance

What ENISA's 2026 State of Cybersecurity in the Union Means for Vulnerability Handling Obligations

2026-11-19

AI-assisted Rapid Research



CSAI | CSA cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

ENISA's 2026 Report on the State of Cybersecurity in the Union, published on 10 November 2026, states that "frontier AI [is] increasingly industrialising vulnerability discovery." It also concludes that "rapid triage, prioritisation, coordinated disclosure and remediation are becoming more critical bottlenecks than discovery alone" [1]. For organisations covered by NIS2 and manufacturers covered by the Cyber Resilience Act (CRA), this moves the compliance risk away from whether a vulnerability handling process exists and toward whether that process can keep working at machine-scale volume. This note complements CSA's earlier summary of the ENISA report [15]. It focuses on one question: how compliance programs built around fixed reporting clocks and documented handling procedures should adapt when the rate of discovery is no longer limited by human researchers.

The first priority is to verify that the CRA reporting clocks will hold under higher report volume. The CRA's 24-hour, 72-hour and 14-day reporting deadlines for actively exploited vulnerabilities have applied since 11 September 2026 [5][8][9]. Higher discovery volume increases the number of reports that must be screened for evidence of exploitation. Because ENISA also finds that AI is shrinking the time from discovery to weaponisation [1], manufacturers may reach the "awareness" threshold more often, and with less time to establish the facts.

Organisations should also treat the remediation bottleneck as a compliance exposure, not only an operational one. According to a separate ENISA survey of entities in the ICT service management sector, 45% took more than three months to apply critical patches to critical systems [1][12]. Once discovery time shrinks "from days to hours" [1], that gap is likely to become harder to defend under NIS2's risk-management obligations. A related concern is automated and AI-generated patching, which should be governed before organisations rely on it. ENISA warns that critical providers "may be forced to use autonomous patching" and that verifying AI-generated patches "may become a major bottleneck" [1], so change-control evidence will need to cover both risks.

Finally, organisations should plan for national coordinated vulnerability disclosure (CVD) policies to change. The EU Action Plan on Cybersecurity and AI invites Member States to "reflect and update their CVD policies and timings" [1][4]. National disclosure expectations may therefore shift, at different speeds across Member States, as governments respond to that invitation.

Background

The State of Cybersecurity in the Union report is required by Article 18 of the NIS2 Directive. ENISA produces it with the European Commission and the NIS Cooperation Group [2][6]. The 2026 edition is the second, and it keeps the Union-level threat assessment at "substantial" [2]. It reports a revised EU Cybersecurity Index with an EU average of 71.53 out of 100 [2]. Because the methodology changed, Help Net Security notes that the index cannot be compared directly with the 2024 edition [12]. The report sets out five directions for the EU ecosystem: effective implementation, crisis operational capacity, critical infrastructure resilience, EU-level supply chain risks and dependencies, and long-term societal resilience [2]. Vulnerability handling falls mainly under the fourth.

The report addresses AI-driven vulnerability discovery explicitly and at length. In its threat landscape chapter, ENISA notes that "model improvements have already been seen accelerating the discovery and exploitation of vulnerabilities." It cites Anthropic's Claude Mythos model, which ENISA dates to the first quarter of 2026, as "a substantial advancement" with dual-use implications [1]. ENISA also cites external data indicating that the time between discovery and weaponisation "is approaching zero." It says frontier AI is having "a tremendous impact on vulnerability management, by narrowing response and patch timings" [1]. In the supply chain chapter, drawing on the Netherlands NCSC's assessment of Mythos [11], ENISA observes that widespread campaigns can weaponise vulnerabilities "within days, and in some cases even hours or minutes, of disclosure," and that frontier AI "will also bring down discovery time from days to hours" [1].

These conclusions build on ENISA's July 2026 publication *ENISA's View on Cybersecurity in the Frontier AI Era* [3]. They also build on the Commission's EU Action Plan on Cybersecurity and Artificial Intelligence, published on 7 July 2026 [4]. The Action Plan encourages organisations to use AI "to detect and address vulnerabilities more quickly." It promotes implementation of NIS2 and the CRA rather than proposing new AI-specific cybersecurity obligations [4]. The regulatory message for compliance teams is that the existing instruments are expected to absorb the AI shift. NIS2's vulnerability handling and disclosure requirements and the CRA's product lifecycle duties remain the obligations that apply. What changes is the operating environment they are tested against.

The CRA's reporting regime has applied since 11 September 2026. Manufacturers of products with digital elements must report actively exploited vulnerabilities and severe incidents through ENISA's CRA Single Reporting Platform (SRP) [8][9]. The CRA's main essential requirements, including the Annex I Part II vulnerability handling requirements, apply from 11 December 2027 [5][8][12]. Organisations are therefore in an interim period. Reporting is mandatory now, but the full lifecycle handling regime that is supposed to sit behind it is not yet enforceable, and the volume of discovery is rising during exactly this period.

Security Analysis

From discovery scarcity to triage scarcity

The central analytical point in ENISA's report is that AI-driven discovery "lowers the marginal cost of identifying flaws and is likely to increase the volume, speed and complexity of vulnerability reports faster than many organisations can triage, validate and remediate them" [1]. In CSA's assessment, most vulnerability handling programs were sized for the discovery rates of the past decade. NIS2 risk-management measures under Article 21(2)(e) require "security in network and information systems acquisition, development and maintenance, including vulnerability handling and disclosure" [6]. For digital infrastructure providers, managed (security) service providers and certain digital providers, Commission Implementing Regulation (EU) 2024/2690 adds more detailed technical and methodological requirements, including on patch management and vulnerability handling [7]. Neither instrument specifies throughput expectations, and in practice compliance with them has typically been demonstrated through documented procedures rather than capacity evidence. When the input rate rises sharply, a process that meets its documented service levels at historical volumes can quietly fall behind without any visible breakdown.

This matters because NIS2 gives supervisors powers to test implementation, not only to review documentation, including on-site inspections, targeted security audits and security scans (Articles 32–33) [6]. ENISA's own data shows how much exposure already exists before any AI-driven increase. As noted above, 45% of surveyed ICT service management entities took more than three months to apply critical patches to critical systems, and 43% conducted no cybersecurity assessments in 2024 [1]. ENISA notes that these failures carry "consequences extending beyond the sector itself to all entities that depend on those providers" [1]. For essential and important entities that rely on managed service providers, this suggests that supplier patch latency should be treated as a NIS2 supply chain risk under Article 21(2)(d), not only as a contractual performance issue. CSA's November 2026 analysis of the disclosure bottleneck reaches a similar conclusion. It argues that the measure that matters is an organisation's capacity to consume patches, not the number of vulnerabilities found [13].

The CRA reporting clock under AI-scale volume

The CRA's reporting timelines are fixed. A manufacturer must submit an early warning within 24 hours of becoming aware of an actively exploited vulnerability, a vulnerability notification within 72 hours, and a final report no later than 14 days after a corrective or mitigating measure becomes available. Severe incidents follow a 24-hour, 72-hour and one-month sequence [5][9]. Practitioner commentary interprets awareness as a "reasonable degree of certainty" that active exploitation is occurring, with reference to

the European Commission's CRA guidance [9]; organisations should confirm this reading against that guidance on Article 14. Manufacturers must also inform impacted users, and where appropriate all users [9].

AI-industrialised discovery puts pressure on this structure in several ways. A higher volume of inbound reports increases the number of candidate vulnerabilities that must be screened for evidence of exploitation. Report volume alone does not trigger the clock, because the CRA obligation attaches to exploitation rather than to discovery. The bridge between the two is ENISA's finding that AI is also accelerating attackers, with the time between discovery and weaponisation "approaching zero" [1]. If exploitation follows disclosure within hours, the period between a report arriving and an exploitation signal appearing also shrinks. The result is that manufacturers may reach the awareness threshold more often while having less time to determine the facts. The 14-day final report is tied to the availability of a corrective measure. If remediation is the bottleneck, as ENISA argues, final reports may stay open longer, and manufacturers may need to rely more on interim mitigations to meet their obligations. CSA's analysis observes that regulatory deadlines of this kind generally assume a fix exists, which the bottleneck increasingly fails to provide [13].

The table below summarises where AI-industrialised discovery is likely to put pressure on specific NIS2 and CRA obligations. It also shows the kind of evidence a compliance program may need to produce.

Obligation	**Source**	**Effect of AI-industrialised discovery**	**Evidence a program should be able to produce**
Vulnerability handling and disclosure as a risk-management measure	NIS2 Art. 21(2)(e) [6]	Intake volume may exceed triage capacity without any visible process failure	Triage throughput and backlog metrics, capacity planning tied to report volume
Supply chain security	NIS2 Art. 21(2)(d) [6]	Supplier patch latency becomes a larger share of exposure	Supplier patch-time commitments, monitoring of supplier remediation
Significant incident reporting (24h / 72h / one month)	NIS2 Art. 23 [6]	Faster exploitation compresses detection-to-notification time	Classification criteria covering AI-accelerated exploitation

Obligation	Source	Effect of AI-industrialised discovery	Evidence a program should be able to produce
Reporting actively exploited vulnerabilities (24h / 72h / 14 days)	CRA Art. 14 [5][9]	The awareness threshold is reached more often and with less time to verify	Documented awareness criteria, SRP submission runbooks, user notification templates
Remediate "without delay," including via security updates	CRA Annex I Part II [5]	Fix development and validation become the constraint	Remediation SLAs by severity, use of compensating controls, records of patch validation
Coordinated vulnerability disclosure policy	CRA Annex I Part II [5]; NIS2 Art. 21(2)(e), with national CVD frameworks under Art. 12 [6]	National CVD timings may be revised under the Action Plan	CVD policy with a defined review cycle and a contact point that can absorb volume

Read together, the rows show a common pattern. Each obligation was already in force or scheduled before the AI shift, but the evidence that demonstrates compliance moves from the existence of a procedure toward measurements of how that procedure performs under load.

Autonomous patching and AI-generated fixes

The most consequential operational passage in the ENISA report for compliance teams may be its discussion of automated remediation. ENISA writes that "critical service providers may be forced to use autonomous patching as organisational approval is simply too slow." It immediately adds that "automated updates could accidentally break systems and cause unwanted downtime" and that "verifying these AI-generated patches may become a major bottleneck" [1]. This sets up a governance tension. NIS2's implementing rules for digital infrastructure providers [7] and common audit frameworks such as ISO/IEC 27001 expect controlled, approved change. The speed of AI-enabled exploitation, however, may make manual approval itself a source of risk.

The two pressures can likely be reconciled through pre-authorised change. Under that approach, an organisation documents in advance which classes of asset, vulnerability severity and patch source may be deployed automatically, which validation gates must pass, and how rollback is triggered. For CRA manufacturers, a parallel issue applies to the fixes they ship. If AI tools help generate security updates, the manufacturer's Annex I obligations, including the requirement to apply effective and regular tests and reviews of product security, would reasonably be read as extending to those updates [5]. Validation of AI-generated code therefore becomes part of the vulnerability handling process rather than a separate engineering concern. ENISA notes that smaller organisations "may find it more difficult to face the challenges than larger entities" [1]. Elsewhere it also flags a risk that small organisations cannot afford access to AI-enabled testing capabilities [1].

Uneven national CVD maturity

The disclosure side of the ecosystem is also changing. ENISA reports that the number of Member States declaring that they have no national CVD policy, or are still defining one, fell from 17 to 7 since 2024. Active implementation, however, has reached only 15 Member States, up from 10 [1]. ENISA warns that where CVD arrangements remain partial, "vulnerabilities may remain unreported, be handled inconsistently or be remediated too slowly" [1]. Organisations that operate across several Member States, or that receive disclosures routed through different national CSIRTs, should therefore expect continued variation in timing and handling expectations. That variation may increase as Member States respond to the Action Plan's invitation to revisit CVD timings [1][4].

At EU level, the infrastructure for absorbing volume is growing. ENISA maintains the European Vulnerability Database (EUVD), which combines CVE data, CSIRT Network information, vendor advisories, known exploited vulnerability catalogues and EPSS scores [1][10]. ENISA has also served as CVE Root for European entities since November 2025 [1]. Under the CRA, vulnerabilities notified through the SRP are made available on the EUVD, with the manufacturer's agreement, once a corrective measure exists [1][5]. Taken together, these developments suggest that the EUVD is likely to become an increasingly authoritative prioritisation source for EU-regulated entities. Organisations that do not currently use the EUVD may want to assess whether its CSIRT Network and CRA-derived data add value to their prioritisation.

Recommendations

The recommendations below follow the analysis in order of urgency. They begin with steps that test current reporting readiness, move to measures that build evidence of capacity over the next two quarters, and close with longer-term positioning for the CRA's full application date.

Immediate Actions

Compliance and product security leaders should first test the CRA reporting process against volume rather than against a single scenario. A tabletop exercise that simulates several concurrent exploitation signals, rather than one, will show whether awareness determinations, SRP submissions and user notifications can be completed within 24 and 72 hours when staff are already busy. Organisations should document the criteria they use to decide that a "reasonable degree of certainty" of exploitation has been reached [9], and check those criteria against the Commission's CRA guidance. A consistent, recorded standard is likely to be easier to defend than ad hoc judgment. NIS2 entities should confirm that their significant-incident classification criteria cover cases where exploitation follows disclosure within hours.

Short-Term Mitigations

Over the next two quarters, organisations should add capacity metrics to their vulnerability handling evidence. Useful measures include inbound report volume, median time to triage, backlog age and remediation time by severity. These give supervisors evidence of how the process performs in practice, not only how it is documented. Organisations should also define a pre-authorised fast path for emergency patching with explicit validation and rollback criteria, which addresses ENISA's autonomous patching concern within a controlled change framework [1]. Contracts and monitoring of managed service providers should include explicit patch-latency commitments for critical systems, given ENISA's ICT service management findings [1]. Organisations should consider adding EUVD data, including its exploitation-likelihood fields, to prioritisation workflows [10].

Strategic Considerations

Over a longer horizon, organisations should plan on the assumption that AI-driven discovery will be more widespread by the time the CRA's full essential requirements apply on 11 December 2027. Programs designed today should be sized for that environment, not for historical volumes. Manufacturers that use AI to generate fixes should build validation of AI-generated code into their secure development lifecycle now, so that Annex I evidence is ready when the requirements apply. Organisations should also monitor

national CVD policy revisions that follow the Action Plan [4], ENISA's further refinement of its frontier AI recommendations [3], and the proposed targeted amendments to NIS2 [1]. Each may change the timing assumptions built into current procedures. Finally, it is worth assessing whether access to AI-enabled testing tools is available across the organisation's supply chain, since ENISA identifies unequal access as a factor that may widen exposure gaps [1].

CSA Resource Alignment

CSA's most directly relevant prior work is *The Disclosure Bottleneck: When AI Discovery Outpaces Patch Consumption* (November 2026) [13]. It reaches a conclusion consistent with ENISA's central finding: AI has made discovery cheap, while validation, triage, fix development and deployment remain constrained by human capacity. It also argues that regulatory reporting clocks, including those of the CRA, assume a fix exists. Its recommendation to measure patch consumption capacity aligns with the capacity metrics recommended above as NIS2 and CRA evidence. CSA's *AI-Accelerated Vulnerability Discovery and the Patch Debt Crisis* [18] examines the same remediation backlog from the perspective of accumulated patch debt and can be read alongside it.

CSA's expedited strategy briefing *The "AI Vulnerability Storm": Building a "Mythos-ready" Security Program* [14] offers an operational playbook for the same Mythos-class capabilities that ENISA and the Netherlands NCSC cite [1][11]. Its guidance on rebuilding triage, prioritisation and remediation workflows for AI-scale input gives practical content that an organisation can document as part of its NIS2 Article 21 vulnerability handling measures. On the disclosure side, CSA's research note *Reforming Coordinated Vulnerability Disclosure for the Autonomous Bug Hunter Era* [17] addresses how CVD policies and timings may need to change under AI-scale discovery, which is the question the Action Plan puts to Member States.

CSA's research note *ENISA's 2026 Compass: AI-Era NIS2 and CRA Compliance* [15] covers the ENISA report's wider findings, including the EU Cybersecurity Index, cross-regime incident taxonomy and supervisory priorities. This note is intended to be read alongside it, with deeper treatment of vulnerability handling obligations specifically.

For control mapping, the CSA AI Controls Matrix (AICM) v1.1 [16] provides the Threat and Vulnerability Management (TVM) and Application and Interface Security (AIS) domains. Organisations can use these to structure evidence for both NIS2 vulnerability handling measures and CRA Annex I Part II requirements. The AICM's supply chain controls also support the managed service provider patch-latency commitments discussed above.

References

[1] ENISA. "2026 Report on the State of Cybersecurity in the Union." European Union Agency for Cybersecurity, November 2026.

[2] ENISA. "ENISA reveals the State of Cybersecurity in the Union." ENISA News, 10 November 2026.

[3] ENISA. "ENISA's View on Cybersecurity in the Frontier AI Era." ENISA, July 2026.

[4] European Commission. "EU Action Plan on Cybersecurity and Artificial Intelligence." Shaping Europe's Digital Future, 7 July 2026.

[5] European Parliament and Council. "Regulation (EU) 2024/2847 on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act)." Official Journal of the European Union, November 2024.

[6] European Parliament and Council. "Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2 Directive)." Official Journal of the European Union, December 2022.

[7] European Commission. "Commission Implementing Regulation (EU) 2024/2690 laying down rules for the application of Directive (EU) 2022/2555 as regards technical and methodological requirements of cybersecurity risk-management measures." Official Journal of the European Union, October 2024.

[8] ENISA. "The CRA Single Reporting Platform is launched." ENISA News, September 2026.

[9] Hogan Lovells. "EU Cyber Resilience Act: Vulnerability and incident reporting obligations now apply and ENISA's Single Reporting Platform is live." Hogan Lovells Publications, September 2026.

[10] ENISA. "European Vulnerability Database (EUVD) – About." ENISA, accessed November 2026.

[11] National Cyber Security Centre Netherlands. "Anthropic's frontiermodel Mythos vraagt om directe actie." NCSC-NL, 2026.

[12] Help Net Security. "EU cyber scores look healthy, but 45% of surveyed IT providers take months to patch." Help Net Security, 11 November 2026.

[13] Cloud Security Alliance AI Safety Initiative. "The Disclosure Bottleneck: When AI Discovery Outpaces Patch Consumption." Cloud Security Alliance, 6 November 2026.

[14] Cloud Security Alliance CISO Community. "The 'AI Vulnerability Storm': Building a 'Mythos-ready' Security Program." Cloud Security Alliance, April 2026.

[15] Cloud Security Alliance AI Safety Initiative. "ENISA's 2026 Compass: AI-Era NIS2 and CRA Compliance." Cloud Security Alliance, 16 November 2026.

[16] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.

[17] Cloud Security Alliance AI Safety Initiative. "Reforming Coordinated Vulnerability Disclosure for the Autonomous Bug Hunter Era." Cloud Security Alliance, 2026.

[18] Cloud Security Alliance AI Safety Initiative. "AI-Accelerated Vulnerability Discovery and the Patch Debt Crisis." Cloud Security Alliance, 2026.