


CRA Article 14 Is Live: 24-Hour Reporting Meets AI-Speed Discovery

What the EU's exploited-vulnerability reporting duty means for manufacturers, open-source stewards, and AI-assisted security programs

2026-10-06

 AI-assisted Rapid Research



CSAI

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

Article 14 of the EU Cyber Resilience Act (Regulation (EU) 2024/2847) began to bind manufacturers on 11 September 2026. From that date, manufacturers of products with digital elements must report actively exploited vulnerabilities and severe incidents through a staged 24-hour, 72-hour, and 14-day process, more than a year before the Act's main security requirements apply on 11 December 2027 [1][2][4][10]. The duty reaches products already on the EU market, so legacy and long-lived devices are in scope today [1][2].

The trigger is exploitation, not discovery. A vulnerability found by an AI-assisted research tool, an internal red team, or a good-faith researcher does not by itself start the clock, and Espressif's reading of the Act excludes good-faith security research and internal testing findings [4]. The clock starts when the manufacturer becomes aware of active exploitation, which Jones Day, summarizing Commission guidance, describes as having "a reasonable degree of certainty" following an initial assessment [1]. CSA's recent work argues that AI-accelerated discovery compresses the time between disclosure and exploitation [6][7]; if that holds, we would expect a higher volume and tempo of events to cross the exploitation threshold.

Organizations should treat the first weeks of operation as a readiness test. The main gaps are likely to be in determining "awareness" quickly, in mapping third-party and open-source components to affected products, and in staffing a 24-hour decision path. The sections below set out the obligation, analyze where AI-speed discovery stresses it, and recommend actions.

Background

The Cyber Resilience Act entered into force in December 2024 and phases in its obligations. The reporting obligations in Article 14 come first. They apply from 11 September 2026, while the essential cybersecurity requirements, conformity assessment, and most remaining duties apply from 11 December 2027 [1][4][5][10]. The phased schedule gives regulators and manufacturers a period of operational experience with incident and vulnerability reporting before full product-design obligations take effect.

Article 14 requires manufacturers to notify two categories of event: actively exploited vulnerabilities in a product, and severe incidents affecting a product's security [1][3]. Reporting follows three stages. An early warning is due within 24 hours of the manufacturer becoming aware. A fuller notification, including an initial assessment and any mitigations, is due within 72 hours. A final report follows within 14 days after a corrective or mitigating measure becomes available for a vulnerability, or within one month for a severe incident [1][2][5]. Manufacturers must also inform affected users and, where appropriate, all users [1][2]. Jones Day notes that manufacturers have no obligation to report retroactively exploitation that was already known before 11 September 2026 [1].

Reports go through the Single Reporting Platform (SRP) operated by ENISA, which routes notifications to the relevant national CSIRT and to ENISA simultaneously [1][5]. ENISA launched the platform as an "initial operating capability" and has said it will extend functionality in the coming months [5]. One Irish government page states that only notifications submitted via the SRP count toward the statutory obligation, with an email fallback available only if ENISA formally declares the platform unavailable [3]. ENISA's launch notice does not describe an API, so readers should check ENISA's current documentation before building automation around the platform [5].

Penalties for non-compliance can reach €15 million or 2.5% of worldwide annual turnover, whichever is higher, and the regulator may also require products to be withdrawn from the market [1][2]. The obligation falls primarily on manufacturers, but ECIJA reports that importers and distributors who market products under their own brand or substantially modify them take on manufacturer duties, and that open-source software stewards have a narrower duty limited to actively exploited vulnerabilities [2]. ENISA's launch notice indicates that the open-source steward reporting obligations begin on 11 December 2027, not now, so the two sources differ and organizations that depend on stewards should confirm timing with counsel [5].

Espressif's guidance for its ESP32 ecosystem illustrates how the duty cascades through supply chains. Espressif monitors and patches its ESP-IDF software and publishes advisories, but downstream product makers remain responsible for the compliance of the complete product. A third-party component that proves to be actively exploited in a finished product triggers the final manufacturer's reporting duty once product impact is confirmed [3][4]. Further duties, such as upstream reporting of fixes to component maintainers and public disclosure once updates ship, arrive with full application in December 2027 [4].

Security Analysis

Defining "awareness" is the central operational question

The 24-hour window is short, but its practical difficulty lies in when it starts. The Act ties the early warning to the manufacturer becoming aware of active exploitation, and Jones Day's summary of the Commission's guidance indicates that awareness is reached with a reasonable degree of certainty following initial assessment, not at first signal [1]. That wording gives manufacturers room to triage, but it does not give them room to delay. A manufacturer that receives a customer report of suspicious activity, a threat-intelligence feed entry listing a component CVE as exploited, or an inclusion of a CVE in a known-exploited catalog will need a documented process for deciding, quickly, whether the evidence is sufficient. We infer that regulators may later examine the timestamps in these decision records, so recording when a signal arrived and how it was evaluated is itself a compliance control.

Exclusions narrow but do not remove the problem. According to Espressif's reading, internal testing findings, good-faith research, and published proof-of-concept code without evidence of real exploitation fall outside the duty [4]. In an AI-assisted environment, however, the boundary between "public proof-of-concept" and "observed exploitation" is likely to blur faster than it did before, which we expect because AI tooling lowers the cost of turning a public flaw into working exploit code. The distinction will therefore depend on evidence from telemetry, customers, and third-party intelligence, not on the mere existence of exploit code.

AI-speed discovery raises the event rate and shortens the response window

CSA's recent work on AI-accelerated vulnerability discovery describes a structural shift in which the bottleneck has moved from finding flaws to remediating them. CSA's patch-debt white paper reports Google Threat Intelligence Group estimates that mean time-to-exploit fell below zero in recent years, meaning exploitation was on average observed before a patch was publicly available, while the same paper puts median remediation of half of internet-facing vulnerabilities at many months [8][6]. We have not independently verified those underlying figures, and readers should consult the cited CSA papers and their original sources. If the pattern holds, Article 14 creates a regulatory clock that runs on the exploitation timeline, not the patch timeline. A manufacturer may have to file a 24-hour early warning for a flaw it cannot fix for weeks.

The early warning is not conditioned on a fix being available. The final report's 14-day clock begins only when a corrective or mitigating measure is available [1][2]. In practice, a manufacturer in a high-volume discovery environment may be filing early warnings and 72-hour notifications for several concurrent

events while the engineering organization is still developing fixes. We expect this to place heavy demands on security-response teams that previously handled such events through coordinated disclosure on a timeline of the researcher's and vendor's choosing [7]. AI-assisted detection and triage could partly offset this by shortening the time needed to establish awareness, although we are not aware of evidence on that effect yet.

Legacy products and deferred maintenance become reportable exposure

Because the duty applies to products already on the market, a manufacturer cannot rely on the December 2027 date to defer preparation [1][2]. Older firmware, discontinued product lines, and embedded devices with long lifecycles can be the subject of an Article 14 report if they are actively exploited, regardless of whether they were designed to the Act's essential requirements. AI-assisted analysis of binaries and open-source dependencies plausibly makes such products more attractive targets, since flaws in code that has not been audited in years can be found at lower cost [6][8]. This suggests manufacturers should inventory the products still in support and know which of them depend on components with active exploitation records.

Dependency mapping determines whether the 24-hour clock is achievable

Most products with digital elements incorporate third-party and open-source components. When one of those is reported as actively exploited, the manufacturer must determine whether its own product is affected and exploitable in its configuration. Without an accurate, queryable software bill of materials linked to deployed product versions, this determination is likely to take days instead of hours. The Espressif example, in which a platform vendor patches and advises but the downstream manufacturer carries the report, shows why component-level traceability is a precondition for timely reporting [4]. Regardless of whether the SRP offers automation, organizations should plan internal workflows that generate report content in advance and keep the submission step simple [5].

AI systems in the product raise scoping and incident questions

For manufacturers shipping products with embedded AI components, such as local models, agent runtimes, or plugin ecosystems, the Article 14 categories will apply to flaws in those components as they do to any other software. A prompt-injection or tool-abuse flaw exploited in the field against a product's agent runtime could plausibly qualify as an actively exploited vulnerability or a severe incident, depending on its effect on the product's security. The regulatory text and early guidance we reviewed do not address such cases specifically, so we treat this as an inference and recommend that manufacturers raise it with counsel and, where possible, with their national CSIRT.

Summary of reporting stages

Stage	Deadline	Applies to	Principal content
Early warning	24 hours from awareness	Actively exploited vulnerabilities; severe incidents	Baseline facts that exploitation or an incident is occurring
Notification	72 hours from awareness	Both	Initial assessment, severity, any temporary mitigations
Final report (vulnerability)	14 days after a corrective or mitigating measure is available	Actively exploited vulnerabilities	Root cause, corrective action taken
Final report (incident)	One month after the 72-hour notification	Severe incidents	Full analysis and measures taken

Sources: [1][2][5]. ECIJA gives the incident final-report deadline as one month [2], and Jones Day anchors it to the 72-hour notification [1]; readers should confirm the exact anchor event against the Regulation [10].

Recommendations

Immediate Actions

Organizations that place products on the EU market should first confirm, in writing, who owns the Article 14 decision and who can submit through the SRP, including out-of-hours coverage. They should register for SRP access ahead of an incident, since onboarding during an active event will consume part of the 24-hour window. They should also define internal criteria for "awareness" and record the timestamps of each signal and decision, so that the reasoning behind a reporting decision can be reconstructed later. Finally, they should identify which in-market products and versions are still supported and which depend on components that appear in known-exploited vulnerability catalogs.

Short-Term Mitigations

Within the next quarter, manufacturers should connect their SBOM and component inventory to deployed product versions so that a component-level exploitation signal can be mapped to affected products in hours. They should prepare report templates for the early warning and 72-hour notification so that staff are completing known fields, not designing a report under time pressure. They should run a tabletop exercise that simulates a 24-hour clock for a third-party component exploited in the wild, including the user-notification step and a decision on how to proceed when no fix exists. Security teams that use AI tools for discovery should also document how those findings are triaged, so that internal research results are separated clearly from exploitation evidence under the Act's exclusions [4].

Strategic Considerations

Over the longer term, manufacturers should plan for sustained notification volume rather than occasional events, in line with CSA's observation that AI-assisted discovery shifts the constraint to remediation capacity [6]. That implies investment in patch engineering, compensating controls that can be deployed before a fix, and supplier contracts that require component vendors to share exploitation intelligence quickly. Manufacturers should also prepare for the December 2027 duties, including upstream reporting to component maintainers and public disclosure after updates ship, because those obligations will interact with the reporting process they are now establishing [4]. Organizations that depend on open-source components should engage with the steward community, since the steward obligations and the practical capacity of maintainers will affect how quickly upstream exploitation information reaches downstream manufacturers [2][5][7]. Finally, we suggest monitoring ENISA's SRP development and any national CSIRT guidance, since the platform is described as an initial capability that will change [5].

CSA Resource Alignment

CSA's research on patch capacity and disclosure supplies the background for this note's central argument. *AI Vulnerability Discovery Is Outpacing Patch Capacity* [6] treats remediation capacity, not discovery, as the limiting factor for defenders, and *AI-Accelerated Vulnerability Discovery and the Patch Debt Crisis* [8] develops the same theme at white-paper length. Neither paper addresses Article 14, so the asymmetry we describe between a 24-hour exploitation clock and a slower patch cycle is our own application of their framing to the new reporting duty. *Project Glasswing and the AI Vulnerability Disclosure Velocity Crisis* [7] argues that the traditional ninety-day coordinated disclosure window was

designed for human-paced research and is strained by AI-paced discovery. In our reading, Article 14 is an early binding mechanism that addresses part of that strain, since the SRP routes notifications to ENISA and national CSIRTs simultaneously.

For readers focused on the Act itself, CSA has published research notes on the Single Reporting Platform [9] and on ENISA's SME maturity model, which exposes a CRA readiness gap [12]; the latter is useful for manufacturers sizing their own preparedness. For control mapping, the threat and vulnerability management and application security domains of the AI Controls Matrix (AICM v1.1) [11] provide a basis for the awareness, triage, and response procedures recommended above.

References

- [1] Jones Day. "[EU Cyber Resilience Act: 24-Hour Reporting Duties Start September 11, 2026](#)." Jones Day Insights, July 2026.
- [2] ECIJA. "[Cyber Resilience Act: Obligations to Report Vulnerabilities and Incidents from September 2026](#)." ECIJA News and Insights, 2026.
- [3] National Cyber Security Centre (Ireland). "[Cyber Resilience Act](#)." NCSC Ireland, 2026.
- [4] Espressif Systems. "[ESP32 CRA Obligations and Deadlines](#)." Espressif Developer Portal, September 2026.
- [5] ENISA. "[The CRA Single Reporting Platform Is Launched](#)." ENISA News, September 2026.
- [6] Cloud Security Alliance. "[AI Vulnerability Discovery Is Outpacing Patch Capacity](#)." CSA Labs, 2026.
- [7] Cloud Security Alliance. "[Project Glasswing and the AI Vulnerability Disclosure Velocity Crisis](#)." CSA Labs, May 2026.
- [8] Cloud Security Alliance. "[AI-Accelerated Vulnerability Discovery and the Patch Debt Crisis](#)." CSA Labs, 2026.
- [9] Cloud Security Alliance. "[CRA Single Reporting Platform \(research note\)](#)." CSA Labs, September 2026.
- [10] European Parliament and Council. "[Regulation \(EU\) 2024/2847 \(Cyber Resilience Act\)](#)." Official Journal of the European Union, November 2024.
- [11] Cloud Security Alliance. "[AI Controls Matrix \(AICM\) v1.1](#)." CSA, 2026.
- [12] Cloud Security Alliance. "[ENISA's SME Maturity Model Exposes a CRA Readiness Gap](#)." CSA Labs, July 2026.