


EU CRA Reporting Begins September 11, 2026

What Manufacturers Need to Operationalize Now

2026-09-04

 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) applies from September 11, 2026. From that date, manufacturers of products with digital elements made available on the EU market must report actively exploited vulnerabilities in their products, and severe incidents having an impact on the security of those products, through ENISA's Single Reporting Platform to the CSIRT designated as coordinator for the relevant Member State [1][3]. Most other CRA obligations – the essential cybersecurity requirements, vulnerability-handling processes and support-period duties – apply from December 11, 2027 [1]. This is a phased application of obligations set out in Article 71, not an accident of sequencing: the Regulation was proposed in 2022 and entered into force in December 2024, and the phase-in was designed to bring reporting online first [1].

The reporting clock is triggered by awareness, and awareness has a defined threshold. Under the European Commission's guidance of July 27, 2026, a manufacturer is considered aware once an initial assessment gives it a reasonable degree of certainty that a vulnerability in its product is being actively exploited, or that a severe incident affecting the security of its product has occurred [2][11]. That initial assessment is expected to be prompt, but it is not a completed forensic investigation, root-cause analysis or remediation. The obligation is to assess quickly and then report within 24 hours of reaching that threshold – not to report every unverified alert, and not to wait for certainty that only a full investigation could provide.

Not every security incident at a manufacturer is reportable. Article 14(5) sets a product-security nexus: an incident is severe where it negatively affects, or is capable of negatively affecting, the product's ability to protect the availability, authenticity, integrity or confidentiality of sensitive or important data or functions, or where it has led or could lead to malicious code being introduced or executed in the product or in a user's network and information systems [1]. A corporate phishing compromise with no effect on a covered product does not become reportable simply because it happened to a manufacturer.

Two operational points are commonly missed. First, Article 14(8) imposes a separate duty to inform users – impacted users, and where appropriate all users – of the vulnerability or incident and, where necessary, of risk mitigation and corrective measures they can deploy [1]. This is a distinct workstream from the CSIRT filing, with its own audience, content and channels. Second, as of the date of this note the Single Reporting Platform is not yet open. ENISA states the platform is scheduled to be operational from September 11, 2026, that its public URL will be published before go-live, and – importantly – advises manufacturers and open-source software stewards to register and initiate validation **only when they**

need to submit a notification [4]. Organizations should therefore not attempt to pre-register, and should direct preparation effort at the things they can actually do now: determining scope, identifying the correct CSIRT, assigning roles, creating personal EU Login accounts, and rehearsing the initial-assessment decision.

Revision Note

This note was updated on September 4, 2026 in response to reader comments on the version published September 3, 2026, and supersedes it. Version 2.0 corrects the described status of the Single Reporting Platform and the associated registration advice, restates the notification limit applicable to non-validated Assigned Representatives, adds the Commission's July 27, 2026 interpretation of the awareness trigger, adds a precise definition of main establishment under Article 14(7) together with the non-EU fallback order, adds coverage of the Article 14(8) user-notification duty, corrects the characterization of the ENISA SME survey population and of the June 11, 2026 conformity-assessment milestone, and reframes the phased application of obligations. Platform-status statements are time-qualified throughout and should be rechecked against ENISA's page before reliance.

What Changes on September 11, 2026

The CRA applies its obligations in stages under Article 71, and only one stage takes effect this month. From September 11, 2026, Article 14's reporting duties apply to manufacturers of products with digital elements. The duty attaches to products already on the market, not only to new releases, so an actively exploited vulnerability in a product shipped years ago is within scope from that date forward [1][3].

The other CRA milestones sit either side of it. Chapter IV, governing the notification of conformity assessment bodies, began to apply on June 11, 2026 [1]. That is a change in the legal status of the rules on notified bodies; it is not the same as a factual claim that conformity assessment bodies started accepting work, and as of late June 2026 no notified bodies had yet been designated in the Commission's NANDO database [10]. The substantive engineering obligations – the essential cybersecurity requirements in Annex I, documented vulnerability handling, and defined support periods – apply from December 11, 2027 [1].

Date	What applies
December 10, 2024	CRA enters into force

Date	What applies
June 11, 2026	Chapter IV, notification of conformity assessment bodies
September 11, 2026	Article 14 reporting obligations
December 11, 2027	Essential cybersecurity requirements and remaining obligations

The fifteen-month interval between reporting and the essential requirements is best read as a deliberate phase-in rather than as evidence that the law demands disclosure before preparation was possible. Attentive manufacturers have had since 2022 to plan. That said, readiness is uneven, and there is credible evidence of difficulty among smaller manufacturers specifically: ENISA's SME CRA survey, fielded in February and March 2026, found that while 66 percent of respondents had heard of the CRA, practical understanding lagged awareness, and the combined domain of incident response and product life-cycle management scored lowest of the five maturity areas at an average of 2.61 out of 5, with 36 percent of microcompanies reporting no incident response plan at all [5]. That survey drew 194 responses, of which 174 were in CRA scope, from 31 countries and geographical groupings – 25 EU Member States alongside Norway, Switzerland and the United Kingdom, plus respondents operating EU-wide and EEA-wide [5]. ENISA states plainly that participation was voluntary and the sample likely reflects a more engaged group of SMEs, so awareness and preparedness across the broader population may be lower [5]. These findings describe surveyed organizations and should not be generalized into an industry-wide readiness verdict.

Penalties give the deadline its weight. Non-compliance with the Article 13 and 14 obligations sits in the top tier of Article 64, attracting administrative fines of up to €15 million or 2.5 percent of total worldwide annual turnover, whichever is higher [1][12].

When the Clock Starts

The 24-hour early warning runs from the point at which a manufacturer becomes aware. The Regulation does not define awareness exhaustively, and the Commission's guidance of July 27, 2026 is the principal interpretive source. That guidance is non-binding – it represents the Commission's reading, and authoritative interpretation ultimately rests with the Court of Justice of the European Union – but it is the clearest available signal of how the threshold will be applied in practice [2].

On the Commission's interpretation, a manufacturer should assess a suspicious event promptly, and is considered aware once that initial assessment produces a **reasonable degree of certainty** that a vulnerability in its product is being actively exploited, or that a severe incident has occurred and compromised the security of its product [2][11]. Awareness may arise from an external report – a customer, a security researcher, a threat-intelligence feed – or from the manufacturer's own monitoring [11].

Two boundaries matter operationally. The threshold is above a raw alert, an unverified allegation or a single anomalous signal: an initial assessment is required before the clock starts. But it is well below a completed investigation. A manufacturer may not defer the early warning until forensic analysis is finished, root cause is established or a fix is designed; the 72-hour, intermediate and final stages exist precisely because the early warning is expected to be filed on incomplete information [1][11]. The practical consequence is that organizations need a named decision-maker with the authority and the criteria to declare reasonable certainty under time pressure, and that decision should be rehearsed before it is needed.

What Is and Is Not Reportable

Article 14 covers two triggers: an actively exploited vulnerability in a product with digital elements, and a severe incident having an impact on the security of such a product [1]. The second is where scoping errors concentrate, because the criterion is the effect on the product, not the fact of a breach at the manufacturer.

Article 14(5) provides the test. An incident is severe where it negatively affects, or is capable of negatively affecting, the ability of the product to protect the availability, authenticity, integrity or confidentiality of sensitive or important data or functions; or where it has led, or is capable of leading, to the introduction or execution of malicious code in the product or in the network and information systems of a user of the product [1].

Applied to common situations, that yields reasonably clear answers. A phishing compromise of a corporate mailbox, or ransomware on an internal finance system, with no path to a covered product, is not a CRA-reportable severe incident – other regimes such as NIS2 or the GDPR may still apply. A compromise of a build system, code-signing key, update-distribution channel or product-support environment is a different matter: those sit directly on the product's integrity and authenticity, and an incident there is capable of negatively affecting the product's protective functions or of introducing malicious code into it. An incident in which malicious code reached a product or reached users through the product falls squarely within the second limb of Article 14(5).

Scope also determines whether a product is covered at all, and the treatment of services deserves care. Stand-alone software-as-a-service is generally outside the CRA and addressed by other instruments. But the Regulation reaches remote data processing solutions where they form part of a product with digital elements. The Article 3(2) definition turns on three conditions: the data processing is carried out at a distance, the product cannot perform one of its functions without it, and the relevant software was designed and developed by the manufacturer or under the manufacturer's responsibility [1][2]. Cloud functionality a manufacturer builds and its product depends on may therefore be in scope even though a general SaaS offering is not, and a third-party service the manufacturer neither designed nor commissioned generally is not [2]. Manufacturers should also confirm the treatment of open-source components and the distinct, lighter regime applying to open-source software stewards, and remember that integrating an open-source component into a commercially supplied product does not transfer the manufacturer's obligations elsewhere [1][2].

Reporting Workflow

Article 14 prescribes a staged sequence, and the final stage differs by trigger. Collapsing the two final-report deadlines into a single generic figure is a common error.

Stage	Deadline	Content
Early warning	Without undue delay, and no later than 24 hours after awareness	Indication of the actively exploited vulnerability or severe incident; where applicable, the Member States in which the product is available
Notification	72 hours after awareness	General information about the product, the general nature of the exploit or vulnerability, and any corrective or mitigating measures
Intermediate report	On request	Status updates, where the coordinating CSIRT asks for them
Final report – actively exploited vulnerability	No later than 14 days after a corrective or mitigating measure is available	Vulnerability description, severity and impact; where available, information on the malicious actor; details of the security update

Stage	Deadline	Content
Final report – severe incident	Within one month of the 72-hour notification	Detailed description of the incident, its severity and impact; type of threat or root cause; applied and ongoing mitigation measures

Source: Regulation (EU) 2024/2847, Article 14 [1][3].

A notification is submitted once through the Single Reporting Platform and routed from there, rather than filed separately in each Member State where the product is available [3]. Manufacturers should prepare the data fields each stage requires in advance, because assembling product identifiers, affected version ranges and availability-by-Member-State during an active incident consumes hours the 24-hour window does not have.

User Notification Under Article 14(8)

Article 14(8) creates an obligation distinct from the CSIRT filing and is easily overlooked because it is not part of the platform workflow. After becoming aware of an actively exploited vulnerability or a severe incident, a manufacturer must inform impacted users of the product – and where appropriate, all users – of the vulnerability or incident, and where necessary of any risk mitigation and corrective measures the users can deploy. Where appropriate, that information should be provided in a structured, machine-readable format that is easily automatically processable [1]. The Regulation's timeliness language is qualitative rather than a fixed hour count; manufacturers should not assume a 24-hour deadline applies here. The consequence of delay is specified: where a manufacturer fails to inform users in a timely manner, the notified CSIRTs designated as coordinators may provide the information to users themselves, where proportionate and necessary to prevent or mitigate impact [1].

Meeting this in practice requires capabilities that are organizationally separate from incident response. A manufacturer needs to know who its impacted users are, which implies a current installed-base or customer-contact capability and the ability to resolve an incident to specific products and version ranges. It needs notification channels that scale and that reach users who may not monitor a security portal. It needs legal, product-security, support and communications functions able to approve external content on a compressed timeline. And it needs the user-facing message to remain consistent with what was filed with the CSIRT while not disclosing exploit detail that would increase risk before a fix is broadly deployed. Organizations that already run coordinated vulnerability disclosure will find much of this familiar; those that do not should treat it as a build.

SRP Status and Access

The following reflects ENISA's guidance as of September 4, 2026 and will change at launch. ENISA's SRP FAQ was itself updated on that date and carries an "under construction" marker, so it should be rechecked directly [4].

ENISA states that the platform is scheduled to be operational from September 11, 2026, and that it will be reachable at a dedicated public URL to be communicated on ENISA's SRP page before the platform goes live [4]. That URL had not been published as of this note, and the list of national CSIRTs designated as coordinators was likewise still outstanding [4][8]. Before launch the platform is undergoing user and security testing, in which various CSIRTs have participated [4]. There is no manufacturer-facing test environment, no rehearsal submission, and no public API for reporting at launch, so the final step of any otherwise-automated workflow will be a person completing a form [4][8].

ENISA's registration advice is explicit and runs contrary to the instinct to get set up early: manufacturers and open-source software stewards are advised to register and initiate the validation process **only when they need to submit a notification**, in order to limit validation workload on CSIRTs [4]. Validation is not a precondition for reporting. It takes place in parallel with the reporting process and does not prevent an Assigned Representative from submitting notifications while it is pending [4]. A non-validated Assigned Representative may submit up to **20 notifications for one manufacturer** before validation becomes mandatory [4].

Access is organized around Assigned Representatives. There can be only one Primary Assigned Representative per manufacturer, and up to 20 Secondary Assigned Representatives [4]. Each must hold an EU Login account, and two-factor authentication is required to access the SRP [4]. EU Login accounts are personal, and ENISA does not recommend registering with a functional mailbox [4]. Creating personal EU Login accounts and deciding who will hold the Primary and Secondary roles are therefore legitimate preparation steps now; they are separate from SRP registration, which should wait.

Selecting the Correct CSIRT

Which CSIRT receives a notification follows from Article 14(7), and the governing concept is not the one most organizations assume. For a manufacturer with an establishment in the Union, the main establishment is the Member State **where the decisions related to the cybersecurity of its products with digital elements are predominantly taken** [1]. That is a question about where product-security decision-making actually sits, and it may resolve to a development centre or product-security subsidiary rather than to the registered office, the corporate headquarters or the largest

commercial operation. Where that Member State cannot be determined, the fallback is the Member State in which the manufacturer has the establishment with the highest number of employees in the Union [1].

For manufacturers with no establishment in the Union, Article 14(7) sets an ordered sequence: the Member State in which the authorised representative is established; failing that, the Member State in which the importer placing the product on the market is established; failing that, the Member State in which the distributor making the product available is established; and failing that, the Member State in which the highest number of users of the product are located [1]. Because this determination fixes the receiving authority for every future notification, it should be made and documented in advance, with legal input, rather than resolved during an incident.

Immediate Readiness Actions

The single most valuable step available before launch is determining, precisely, which of an organization's offerings are products with digital elements within the meaning of the Regulation, including any remote data processing solutions that form part of them. Scope determines everything downstream, and it is a legal and product question rather than a security one, so it needs legal and product-management involvement rather than being delegated to a security team alone.

With scope settled, map where cybersecurity decisions for those products are predominantly taken, apply Article 14(7) to identify the main establishment and therefore the coordinating CSIRT, and record the reasoning. Identify who will serve as Primary Assigned Representative and which colleagues will be Secondary Assigned Representatives, and have those individuals create personal EU Login accounts with two-factor authentication – while following ENISA's advice not to initiate SRP registration until a notification actually needs to be submitted [4].

The remaining preparation is procedural. Name the person or role that performs the initial assessment and the person who declares reasonable certainty, and give them written criteria; this is a judgement call under time pressure and it should not be improvised. Document the escalation path that carries a finding from detection to a filed early warning inside 24 hours, and from there to the 72-hour notification. Pre-build the data set each stage requires. Write down the boundary between a product-affecting incident and an unrelated corporate IT incident, so that triage does not have to relitigate the question mid-event. Stand up, or extend, the Article 14(8) user-notification process, including the installed-base data it depends on. Then run two tabletop exercises – one on an actively exploited vulnerability and one on a product-impacting severe incident – because the two follow different final-report tracks. Organizations

already subject to NIS2 should map the two regimes' triggers, timeframes and recipients side by side so one event does not generate conflicting internal deadlines [9]. Finally, monitor ENISA's SRP page for the launch URL, the coordinator CSIRT list and revised instructions [4].

Longer-Term CRA Preparation

The period to December 2027 is a planning window for the obligations that reporting will eventually sit on top of. A current product inventory with accurate version and support-status data underpins both the reporting duty and the later requirements. Software bill of materials generation, a documented vulnerability-handling process, secure-by-default configuration and cryptographically verifiable updates are all Annex I concerns that take longer to build than to describe. Support-period determination and disclosure will need product-management ownership, not only engineering. Enterprises that depend on smaller suppliers in CRA-regulated categories may reasonably fold CRA readiness into third-party risk review, given the survey evidence of weaker capability among SMEs [5] – while keeping in mind that those findings describe surveyed organizations rather than the whole supplier population.

CSA Resource Alignment

CSA's closest companion analysis is [ENISA's SME Maturity Model Exposes a CRA Readiness Gap](#), published July 15, 2026, which works through ENISA's five-domain maturity model against CRA obligations and covers the engineering and governance dimensions this reporting-focused note does not [7]. For the internal mechanics of triaging and escalating an event and exchanging information about it – the capability Article 14 exercises under deadline – CSA's [Best Practices for Cyber Incident Exchange](#) provides a process framework that predates the CRA but maps onto the workflow manufacturers now need to formalize. Organizations placing CRA work alongside their other EU obligations may find CSA's overview of [European compliance across NIS2, CRA and DORA](#) useful for separating the regimes' distinct triggers and recipients. Where an organization already maps its control environment to the [Cloud Controls Matrix](#), the threat and vulnerability management and incident response domains supply the documented control evidence that a defensible 24-hour reporting decision rests on.

References

- [1] European Union. "[Regulation \(EU\) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements \(Cyber Resilience Act\)](#)." Official Journal of the European Union, 2024. See in particular Articles 3, 13, 14, 64 and 71.
- [2] European Commission. "[Commission publishes new guidance to support timely Cyber Resilience Act implementation](#)" – Commission Decision C(2026) 5252 final and its Annex, Commission guidance on the application of the Cyber Resilience Act. Shaping Europe's Digital Future, July 27, 2026. Non-binding guidance.
- [3] European Commission. "[Cyber Resilience Act – Reporting obligations](#)." Shaping Europe's Digital Future, 2026.
- [4] ENISA. "[Single Reporting Platform \(SRP\) – Frequently Asked Questions](#)." European Union Agency for Cybersecurity. Page updated September 4, 2026; consulted September 4, 2026.
- [5] ENISA. "[SME CRA Survey Report](#)." European Union Agency for Cybersecurity, June 2026.
- [6] ENISA. "[CRA SRP guidance – AR user registration](#)." European Union Agency for Cybersecurity, 2026.
- [7] Cloud Security Alliance. "[ENISA's SME Maturity Model Exposes a CRA Readiness Gap](#)." CSA AI Safety Initiative, July 15, 2026.
- [8] ENISA. "[How to Report Incidents and Actively Exploited Vulnerabilities via the CRA Single Reporting Platform](#)." Factsheet v1.0, July 2026.
- [9] Cloud Security Alliance. "[An Update on European Compliance: NIS2, CRA, DORA](#)." CSA Blog, September 18, 2025.
- [10] cyberresilienceact.eu. "[On 11 June 2026, the CRA's Rules on Notified Bodies Started to Apply, but None Are Designated Yet](#)." June 2026.
- [11] Freshfields. "[Cyber Resilience Act reporting obligations take effect on 11 September 2026](#)." Technology Quotient, 2026.
- [12] Freshfields. "[Decoding the Cyber Resilience Act – Part 3: Managing CRA Risk in Practice](#)." Technology Quotient, 2026.