Kunnus

WHITEPAPER · EU CYBER RESILIENCE ACT

Article 14 Reporting Obligations: What Must Be in Place by 11 September 2026

Actively exploited vulnerabilities and severe incidents: notified within 24 hours, detailed within 72 hours, closed out with a final report. This whitepaper turns every obligation into a concrete action, with deadlines, a process map and a checklist.

- Current as of the interpretative practice of July 2026

PUBLISHED BY

Kunnus · Think Ahead Technologies GmbH

Version

July 2026

CONTENTS

In this whitepaper

AT A GLANCE

11 Sep 2026
the reporting obligation applies, including to products already on the market today

24 h / 72 h
early warning and detailed notification, in calendar hours, weekends included

€15m
or 2.5% of worldwide turnover: the penalty range for Article 14 infringements

01

WHY NOW

The first hard CRA deadline

On 11 September 2026, Article 14 of Regulation (EU) 2024/2847, the EU Cyber Resilience Act, becomes binding. From that day, manufacturers of products with digital elements must notify actively exploited vulnerabilities and severe incidents to the coordinating CSIRT of their Member State and to ENISA. The remaining manufacturer obligations, among them vulnerability handling, security updates and CE conformity, only follow 15 months later.

The text of the regulation leaves interpretative room in three places, and those three are where it is decided in practice whether a notification goes out on time: from when a manufacturer is deemed to have "become aware", what applies to existing and end-of-life products, and how vulnerabilities in third-party components are to be treated. This whitepaper answers those questions as the Commission's current interpretative practice stands.

THE CRA ROADMAP

10 Dec 2024
CRA enters into force

27 Jul 2026
Commission guidance on the CRA

11 Sep 2026
Art. 14 reporting obligation applies to all products within the scope of the CRA

11 Dec 2027
Full applicability: vulnerability handling, Annex I, CE obligations

PENALTY RANGE

Infringements of the reporting obligation sit in the highest penalty tier of the CRA: up to **€15 million or 2.5% of worldwide annual turnover**, whichever is higher.

REACH

The reporting obligation **also applies to existing products** (placed on the market before 11 Dec 2027) and, unlike vulnerability handling, **continues after the end of the support period**.

02

TWO TRIGGERS, TWO REPORTING CHAINS

What Article 14 requires

ART. 14(1)–(2)

Actively exploited vulnerabilities

You establish that a vulnerability in your product is **actually being exploited** by attackers. A published CVE alone does not trigger the deadline. Only your awareness of active exploitation starts the clock.

ART. 14(3)–(5)

Severe incidents

An incident that affects the product's ability to protect the **availability, authenticity, integrity or confidentiality** of sensitive data and functions. Likewise an incident that has led or could lead to **malicious code** in the product or a user's system.

> You notify both events **simultaneously** to the CSIRT designated as coordinator in your Member State and to ENISA, via ENISA's Single Reporting Platform under Article 16. It is due to be operational by 11 September 2026, with a testing period beforehand. Only in exceptional cases do you notify by the route your CSIRT publishes — that does not extend the deadline.

What Article 14 does not require

Many manufacturers significantly overestimate the scope of the reporting obligation. The following are not reportable:

- ✕ **Vulnerability scanner findings.** A CVE in a dependency is a matter for your internal vulnerability management.
- ✕ **Reports from security researchers.** Handle and fix, yes; reportable only once active exploitation is established.
- ✕ **Published CVEs without exploitation.** Even at CVSS 9.8, as long as no active exploitation of your product is known.
- ✕ **Pentest and audit findings.** Vulnerabilities you find yourself remain an internal matter.

One special case deserves attention: if a researcher's report shows that something has already happened, such as injected malicious code or a compromised build pipeline, it constitutes a severe incident under Art. 14(5)(b). That is reportable even without established active exploitation. Assess incoming reports against both triggers, always.

03

THE DECISION TREE

Reportable or not?

A lead comes in – researcher, customer, authority, media, your own team

Immediate initial assessment – against both triggers, timestamped

NEITHER TRIGGER MET

No reporting obligation. Handle as a normal vulnerability and document the decision. Voluntary notification under Art. 15 remains open.

ACTIVE EXPLOITATION VERIFIED

Vulnerability chain: 24 h early warning, 72 h detailed notification, final report 14 days after a fix is available.

SEVERE INCIDENT (ART. 14(5)) VERIFIED

Incident chain: 24 h early warning, 72 h detailed notification, final report within one month.

THE MOST COMMON BORDERLINE CASE

Third-party components: an actively exploited vulnerability originating from an integrated component is reportable once it is contained *in your product*. If the vulnerability **cannot be exploited in your product**, for instance because the vulnerable code is not reachable, or **has not been exploited there**, you have no Art. 14 reporting obligation — the Commission FAQ says so explicitly. Document that assessment; the burden of proof sits with you. What remains: vulnerability handling under Annex I Part II, upstream reporting under Art. 13(6) and the option of voluntary notification under Art. 15.

The notification grows with the analysis

Nobody expects finished forensics after 24 hours. The early warning deliberately contains limited information, and the notification is updated progressively as the internal investigation advances.

Voluntary notification under Art. 15

Vulnerabilities and incidents below the threshold can be notified voluntarily to the CSIRT and ENISA at any time, for instance where warning other manufacturers seems sensible. No obligation arises from doing so.

04

THE REPORTING DEADLINES

Three stages at a glance

0 → 24 h

Early warning

Product, title; for vulnerabilities the Member States concerned, for incidents whether malicious action is suspected

→ 72 h

Detailed notification

Nature of the vulnerability or incident, initial assessment, measures taken, workarounds for users

Develop the fix

No statutory deadline

Document when the corrective or mitigating measure becomes available, because it starts the final deadline

14 d / 1 m

Final report

Vulnerability: 14 days after the measure is available · Incident: 1 month after the 72-hour notification

STAGE	ACTIVELY EXPLOITED VULNERABILITY	SEVERE INCIDENT
Early warning	within 24 hours of becoming aware	within 24 hours of becoming aware
Detailed notification	within 72 hours of becoming aware	within 72 hours of becoming aware
Final report	no later than 14 days after a corrective or mitigating measure is available	within one month of the 72-hour notification

Calendar hours, not office hours. The 24 and 72 hours keep running over weekends and public holidays. Verify on Friday evening, and you notify by Saturday evening at the latest. Plan your on-call cover accordingly.

Status updates on request (Art. 14(6))

The CSIRT may request status updates at any time. Keep a case timeline from the first notification and such requests become an export rather than a reconstruction.

Keep an eye on formats (Art. 14(9)–(10))

The Commission specifies notification formats and procedures by implementing act. The mandatory fields per stage are listed in the ENISA FAQ on the Single Reporting Platform; check your templates against it quarterly.

05

BECOMING AWARE

When the clock really starts

Every deadline starts with your becoming aware, and that is not the moment a report lands in your inbox. Awareness presupposes that, after an initial assessment, you can conclude with **a reasonable degree of certainty** that active exploitation is occurring or that a severe incident has compromised the security of your product. This yardstick is no CRA special case: it matches the concept of awareness in the NIS2 Implementing Regulation (EU) 2024/2690 and established practice on personal data breach notification under the GDPR. Whoever has processes there can reuse them here.

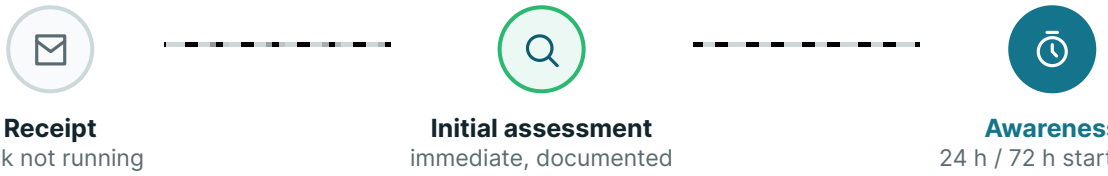

Awareness means verification with a reasonable degree of certainty, not the inbox. The assessment time in between does not count towards the deadline, but it must start promptly and be documented.

The initial assessment cannot be deferred. What is required is that you assess every suspicious event **immediately**, whether it comes from a customer, a security researcher, an authority or the media. Dragging out verification does not postpone the deadline; it invites the accusation of having become aware too late. So record when the report arrived, when the assessment started and when the result was established.

No retroactivity, with one important exception. Exploitation you were already aware of before 11 September 2026 does not have to be reported retroactively. But if you only knew of the vulnerability and not of its exploitation, and active exploitation occurs or comes to your attention after the deadline, it is fully reportable.

06

BORDERLINE CASES

Four cases that are contested in practice

Existing products and end of support

From 11 Sep 2026 the reporting obligation applies to **all** products within the scope of the CRA, including those placed on the market before 11 Dec 2027. Nor does it end with support: unlike vulnerability handling, it continues after the support period ends. Only the Annex I Part II obligations fall away then.

Upstream reporting (Art. 13(6))

You report vulnerabilities in integrated components to their manufacturer or maintainer, but only for the **version you integrate** and only for vulnerabilities of the component itself, not of the integration. The obligation lapses where the maintainer is demonstrably already aware (check CVE databases, advisories and issue trackers first to avoid duplicates) or where there is no maintainer any more.

Informing users under Art. 14(8)

Impacted users must be informed, but not necessarily publicly. Details may be **limited to those affected, on a risk basis**, particularly for products in sensitive environments where public technical detail would facilitate further attacks. After remediation, broader disclosure may be appropriate. Fixed vulnerabilities must be disclosed with the update under Annex I Part II point (4).

Open-source stewards (Art. 24)

How far Art. 14 applies to FOSS stewards depends on the type of support they provide. Those providing **IT infrastructure** report severe incidents affecting that infrastructure. Those contributing **engineering resources** report actively exploited vulnerabilities and inform users where appropriate. The obligation only bites on awareness of active exploitation, typically via reports from manufacturers who have integrated the component.

> None of these four cases tightens the reporting obligation; all four make it plannable. Set up initial assessment, documentation and responsibilities properly and you have settled them in your favour.

07

THE PROCESS MAP

From lead to completed report

Clock not running

1

Receipt

A lead arrives, whether from a researcher, a customer, your own team, an authority or the media, and lands in the defined intake channel. Document receipt with a timestamp.

Duty to assess

2

Initial assessment, immediately

Assess against both triggers. If neither is met, there is no reporting obligation; the decision is documented. If verified, the moment of verification is your awareness. The clock starts now.

Hour 0–24

3

Early warning

Via the Single Reporting Platform to the CSIRT and ENISA: product, title, Member States concerned or suspected malicious action. With a prepared template this is a form-filling exercise. Immediate technical measures run in parallel.

Hour 24–72

4

Detailed notification

Follow up with the analysis: nature of the vulnerability or incident, initial assessment, measures taken, workarounds. User notification under Art. 14(8) runs in parallel and begins with awareness, not here: inform known impacted customers without delay, and the broader user base via an advisory as soon as that adds no extra risk.

No deadline

5

Develop the fix

Develop the patch or mitigation. Article 14 sets no deadline for this. The moment the measure becomes available starts the final deadline, though, so it belongs in the case file.

14 d / 1 month

6

Final report

For vulnerabilities no later than 14 days after the measure is available, for incidents within one month of the 72-hour notification. The case timeline and the security advisory provide the content, the advisory machine-readable as CSAF where useful.

Throughout, phases 1 to 6: document. Every step timestamped in the case timeline. CSIRT requests for status updates (Art. 14(6)) then become an export, and you can prove at any time what you knew and did, and when.

08

A WORKED EXAMPLE

The real thing at an IoT gateway manufacturer

A fictitious but realistic case. This is how the chain plays out once the preparation from chapter 09 is done:

Tue, 14:32	A security researcher reports a vulnerability in the remote-maintenance module via the intake channel. He includes a proof of concept, but no indication that the flaw is being exploited. Receipt logged with a timestamp.
Tue, 15:05	The initial assessment begins: reproduce the vulnerability, check telemetry and customer logs against the attack pattern. Assessment start documented.
Tue, 19:20	**Awareness.** The access pattern matches in two customers' logs; active exploitation is confirmed. The clock is running, the early warning is due by Wednesday, 19:20.
Wed, 08:45	The early warning goes out via the Single Reporting Platform to the CSIRT and ENISA, with product, title and Member States taken from the prepared template. Ten hours of buffer left unused, because the template was ready.
Thu, 14:10	The detailed notification goes out, comfortably within the 72-hour deadline: nature of the vulnerability, initial assessment, a workaround (disable the remote-maintenance module). Known impacted customers are informed directly in parallel.
Nine days later	The patch is tested and available. The moment is documented; the 14-day deadline for the final report starts now.
Three days after that	Final report submitted, security advisory published as CSAF. The case is closed, and the complete timestamped timeline is on file.

> Between receipt and the start of the deadline lie **just under five hours of documented assessment time** — the report did not yet evidence exploitation. The clearer a report already makes it, the shorter that latitude: assessment time does not extend the deadline, it has to start without delay and be provable.

09

PREPARATION

Four steps to 11 September

- [ ] **1 · Set up the intake channel and initial assessment**
  One defined intake channel for all leads plus a classification matrix under Art. 14(5) with examples from your own product world. That way the team decides in minutes whether a case is reportable. A public intake channel only becomes formally mandatory from December 2027 (Art. 13(17)); without one, though, you make life unnecessarily hard for yourself and end up hunting for leads across scattered inboxes when it matters.

- [ ] **2 · Assign responsibility**
  A reporting owner plus a deputy named, the competent CSIRT documented. What counts is the Member State of the main establishment, meaning where the cybersecurity decisions are predominantly taken. Set up access to the Single Reporting Platform and trial it during the testing period before the deadline.

- [ ] **3 · Build notification templates**
  One template each for the 24-hour early warning, the 72-hour notification, the final report and the user notification. Which fields are mandatory at each stage is listed in the ENISA FAQ on the Single Reporting Platform.

- [ ] **4 · Run it through once**
  A tabletop run-through with a fictitious vulnerability shows where the chain snags, while it still costs nothing.

For most manufacturers the effort amounts to a few person-days. The 24-hour deadline only becomes a problem when processes, responsibilities and templates have to be hunted down mid-incident.

A

APPENDIX · WORKING AIDS

Classification and responsibilities

CLASSIFICATION MATRIX — SIX TYPICAL CASES

SITUATION	CLASSIFICATION	CONSEQUENCE
Attackers demonstrably exploit an RCE flaw in the remote-maintenance module	**Actively exploited vulnerability**	Vulnerability chain, 24 h from awareness
Ransomware encrypts device data at several customers	**Severe incident (Art. 14(5)(a))**	Incident chain, 24 h from awareness
Malicious code discovered in a signed update package, no attack known yet	**Severe incident (Art. 14(5)(b))**	Reportable even without established exploitation
CVE 9.8 in a library you use, no exploit against your product known	**Neither trigger**	Handle internally, upstream under Art. 13(6) where relevant
A pentest finds an auth bypass in your own firmware	**Neither trigger**	Internal matter, fix via the vulnerability process
DDoS takes down your company website, the product is unaffected	**Not a product incident**	Not Art. 14; check NIS2 reporting duties instead

ROLES IN A LIVE CASE — NAME THEM IN ADVANCE, DO NOT IMPROVISE

ROLE	DECIDES ON	AVAILABILITY
Reporting owner	Establishing awareness, approving and sending the notification	24/7, with a deputy
Technical analysis	Reproduction, evidence of exploitation, workaround and patch	On call
Customer communication	User notification under Art. 14(8), advisory and messaging	Business hours
Management	Escalation, recall or sales stop, external support	Escalation path

A

APPENDIX · WORKING AIDS

Notification content and the run-through

NOTIFICATION CONTENT PER STAGE — ADOPT AS TEMPLATE FIELDS

Early warning (24 h)

Manufacturer and contact point · affected product with version · vulnerability or incident · for vulnerabilities: EU Member States concerned · for incidents: suspected malicious action, cross-border impact

Detailed notification (72 h)

Nature of the vulnerability or incident · initial assessment of severity and impact · CVE identifier where one exists · corrective and mitigating measures taken · workarounds for users

Final report, vulnerability (14 days after fix)

Description with severity and impact · information on the attacker, where known · available update or corrective measure

Final report, incident (1 month)

Detailed description with severity and impact · type of threat and likely root cause · remedial measures taken and ongoing

A 60-MINUTE TABLETOP EXERCISE — THE RUN-THROUGH AGENDA

0–10 min	**Inject the scenario.** A fictitious researcher report with a log extract goes to the real intake channel. Who notices it, and how long does it take?
10–25 min	**Initial assessment against the matrix.** The team decides using the classification matrix: active exploitation, severe incident, or neither trigger? Record the decision and the timestamp.
25–40 min	**Fill in the early warning.** Populate the template with real product data, without sending it. Every field somebody has to go and look up is a finding.
40–50 min	**Draft the user notification.** Who are the impacted customers, through which channel do you reach them, and who approves the wording?
50–60 min	**Log the gaps.** Every finding becomes a task with an owner and a date. Only then was the exercise an exercise rather than a discussion.

FREQUENTLY ASKED QUESTIONS

Answered briefly

Do I have to report every CVE from my scanner?

No. Only actively exploited vulnerabilities and severe incidents under Art. 14(5) are reportable. Scanner findings, CVEs without exploitation and pentest results belong in your internal vulnerability management.

When does the 24-hour deadline start?

With verified awareness, not with the receipt of a report. The initial assessment must begin immediately, however, and be documented.

Who do I report to?

Simultaneously to the coordinating CSIRT of the Member State of your main establishment and to ENISA, via the Single Reporting Platform under Art. 16.

Does this apply to old, no-longer-supported products too?

Yes. The reporting obligation applies to all products within the scope of the CRA and continues after the end of the support period. Only vulnerability handling ends with support.

We have no establishment in the EU. Which CSIRT is competent?

Art. 14(7) provides a cascade: the Member State of the authorised representative, failing that the importer, failing that the distributor, failing that the state where the users are. Determine it once, document it, done.

Is an SBOM required for the notification?

Not for Art. 14; it only becomes mandatory from December 2027 via Annex I. Without one, though, you will struggle to establish within the deadlines which products a component vulnerability affects.

YOUR CONTACT

Maximilian Heck
Head of Regulatory & Security

maximilian.heck@think-ahead.tech
+49 176 72001773 · linkedin.com/in/maximilian-heck

Sources: Regulation (EU) 2024/2847 (Cyber Resilience Act), Arts. 13–16, 24, 64, 69, 71 · Commission guidance on the application of the CRA, July 2026 · Implementing Regulation (EU) 2024/2690 (NIS2) · ENISA FAQ on the Single Reporting Platform. Only the wording published in the Official Journal of the EU is legally binding. This whitepaper does not constitute legal advice.