

The Cyber Resilience Act's first deadline hits in September

The Cyber Resilience Act's first deadline hits in September

□ Bernard Sfez - 2026-08-11 15:34



On 11 September 2026, the Cyber Resilience Act starts requiring "manufacturers" to report any actively exploited vulnerability in their products within 24 hours. A full notification follows within 72 hours, then a final report within fourteen days. The obligation is not limited to future products: it covers the installed base, including a device shipped five years ago and still running at a customer site.

Which leaves the real question: who counts as a manufacturer? The word covers far more people than you would think. Whoever sells a plugin or a module. Whoever assembles a box with embedded software. Whoever ships an image or an appliance under their own brand. Hosting providers, integrators and end users do not qualify, and non-commercial open source falls outside the scope entirely. Here is the test that settles it, the situations where it applies, and what can still be done before the deadline.

The Cyber Resilience Act, Regulation (EU) 2024/2847, has been in force since December 2024 and has spent most of the two years since as a conference topic. On 11 September 2026 it becomes enforceable. Not in full: the reporting obligation goes live first, and it comes with deadlines few organisations can currently meet.

The date is widely misunderstood, for two reasons. First, people confuse it with the full application of the regulation, which comes in December 2027. Second, most assume that "manufacturer of a digital product" describes somebody else: a connected device maker, a software vendor with a hundred employees. The regulation does not see it that way.

Some good news for anyone starting now: on 27 July 2026 the Commission approved its official guidance on applying the regulation (Communication C(2026) 5252). It is not binding, since only the Court of Justice can settle a point of law, but market surveillance authorities will work from it. It contains 67 practical examples and pays explicit attention to micro-enterprises and SMEs. It is the document that was missing to answer the questions small organisations actually have.

EU REGULATION

CYBER RESILIENCE ACT

11 SEPTEMBER 2026



What changes on 11 September, and what does not

What changes:

From that date, any manufacturer who becomes aware of an actively exploited vulnerability in one of its products, or of a severe incident affecting its security, must issue an early warning within 24 hours. It goes to the single reporting platform operated by ENISA, through the entry point of the coordinating CSIRT of the country where you have your main establishment, which is CERT-FR for France. A full notification follows within 72 hours. Then a final report: at the latest 14 days after a fix becomes available for a vulnerability, or one month after the initial notification for a severe incident.

Twenty-four hours means calendar hours. A flaw found under exploitation on a Friday evening in August does not wait until Monday.

Worth being precise here: "actively exploited" is not the same as "published". It requires reliable evidence that a malicious actor has exploited the flaw on a system without authorisation. Public disclosure, a proof of concept or a researcher's demonstration do not, on their own, start the clock.

What does not change:

The design requirements, meaning security by default, the software bill of materials (SBOM), the absence of known exploitable vulnerabilities at the time of placing on the market, and CE marking, only come into play in December 2027, and only for products placed on the market after that date. Your existing catalogue is not caught retroactively.

The reporting obligation, on the other hand, does apply to the installed base. An actively exploited flaw in a product you shipped in 2022 must be reported within 24 hours as of 11 September 2026. That is the distinction everyone gets wrong.



Penalties for failing to meet this obligation reach 15 million euros or 2.5% of worldwide annual turnover. Nobody expects that ceiling to be used against a small business, but the obligation itself does not scale with your size.

The three-question test

You are a manufacturer under the CRA if you answer yes to the first two questions, or to the third one alone.

1. Do you supply a digital product under your own name or brand? A board, a sensor, a controller, a connected box, the firmware running inside it. But also a plugin, a theme, a module sold separately, a Docker image, a preconfigured appliance. Neither the size of the product nor the size of your company matters. Putting your brand on somebody else's product is enough: you become its manufacturer.

2. Is it in the course of a commercial activity? Open source software developed and distributed outside any commercial purpose is out of scope. "Commercial" does not mean "paid for": free software whose support, hosting or pro edition you monetise is in scope.

3. Or: do you substantially modify an existing product before supplying it? A fork you maintain and redistribute commercially makes you the manufacturer of the result, even with nobody's brand on it.

On that third point, the July guidance settles something everyone had been waiting for: a security update is not, in principle, a substantial modification. Its purpose is to reduce risk. The test is whether the product's risk profile changes, not how large the change is. Fixing a flaw does not throw you back into a fresh conformity assessment.

If you answered yes, the rest of this article is about you. And if you would rather have an outside opinion than call it alone, we give one in an hour.

Concrete situations

You design or assemble equipment with embedded software. A sensor, a controller, a kiosk, a medical device. You are a manufacturer, no argument. And here is the part that hurts: the obligation covers the installed base. A unit shipped in 2021, still in service at a customer site, whose firmware includes a library with a flaw now under exploitation: report within 24 hours.

You sell a plugin, a theme or a module. You are a manufacturer. This is the case that gets overlooked most often.

You distribute a Docker image, an appliance or a preconfigured distribution under your brand. You are a manufacturer, even if every piece of it is third-party software, and including for flaws in the components you did not write.

You run an open source project you do not sell, but with structured commercial backing. You most likely fall under the open source software steward status, created by the regulation for exactly

this case. Lighter regime: no CE marking, no formal conformity assessment. But not an empty one. It requires a written cybersecurity policy, cooperation with market surveillance authorities, and, from 11 September 2026, reporting of actively exploited vulnerabilities and of severe incidents affecting the systems you make available for the project's development.

You resell or install a third-party product without touching it. You are a distributor, not a manufacturer. The regulation still gives you duties, lighter ones: do not make available a product you know to be non-compliant, and inform both the manufacturer and the market surveillance authority if you discover a flaw. Checking the CE marking will be added to that list, but only from December 2027. Before then there is nothing to check.

You install and maintain a CMS (Tiki, WordPress), an ERP or a Nextcloud for a client. You are supplying a service, not placing a product on the market. You are not a manufacturer. Neither are your clients: they are users. The same goes for hosting, managed services and monitoring sold on a monthly basis. One caveat: the CRA not applying does not mean nothing applies. Managed service providers may well fall under NIS2, with its own thresholds and its own obligations. That is a different subject, not the absence of one.

You build custom work: a theme, a template, a line-of-business application. The July guidance finally gives a usable criterion, based on where the software runs. What executes on the user's device (a downloaded application, a browser extension, a locally installed client) is a product. A web application used only through a browser, or a site that presents information, generally is not, unless it qualifies as remote data processing necessary for a product to function. A theme built for a Tiki or a WordPress accessed through a browser therefore falls outside. A plugin distributed and installed at the client's end falls inside. And simply publishing source code in a public repository does not generally amount to placing a product on the market.

If you are a manufacturer: the minimum before the 11th

Five things. None of them needs a budget. All of them need a decision.

A public reporting channel. A `security@` address that genuinely exists, a `security.txt` file, a coordinated disclosure page. Without it, you will learn that your flaw is being exploited from an angry customer. This is not optional: a coordinated vulnerability disclosure policy is a requirement of the regulation, written down, published and applied.

Knowing who decides at 2 in the morning. You cannot register in advance: the single reporting platform only goes live on 11 September 2026, and manufacturers register on it from that point. What you can prepare is the decision. Who picks up the phone, who is allowed to declare, and with what information at hand.

An on-call arrangement, however informal. The 24-hour clock does not know about weekends.

An inventory of your dependencies. The SBOM is not required before December 2027, but without one you will not know whether today's CVE concerns you. For a software vendor, a composition analysis tool produces a first list in a day. For a hardware manufacturer it is a different project: tracing the firmware subcontracting chain, knowing which version runs on which batch, and sometimes discovering that nothing deployed can be updated remotely at all. Count in months, not in afternoons, and start now.

Some way of knowing that a flaw is actively exploited. CVE watch, the KEV catalogue, CERT-FR advisories, and above all logs and monitoring that will show you an exploitation attempt in progress. This is where most organisations are blind.

One written procedure page covers all five. Write it before September.

A CVE watch tied to your inventory is not a one-off project, it is a daily routine. That is the kind of thing you hand off: we run it for our clients.

A sixth item, due by December 2027 but decided now: the support period. Five years is a floor, not a default. It has to reflect how long the product can reasonably be expected to be in use, and it has to be stated to the customer at the time of purchase, at least as a month and a year. If you sell hardware that gets installed for ten years, saying five will not hold.

If you are not a manufacturer: it still concerns you

The obligation creates an upstream flow of information. Your suppliers will have to declare their exploited flaws, which makes their diligence verifiable, and therefore contractual.

Four points to add to your contracts and your supplier scorecards: how quickly the supplier informs you (not ENISA, you); the support period stated for the product; whether a disclosure channel is published; and delivery of an SBOM with the product. A supplier who cannot answer those four questions in September 2026 will cost you dearly in December 2027.

What this says about the rest

The CRA formalises an obligation that already existed in practice: know what is running, know when it breaks, know who to tell. Organisations with an inventory, monitoring and an incident procedure will treat 11 September as paperwork. The others will find out that compliance is not a document to produce but an operational capability to build, and that three weeks is not enough to build it.

If you are reading this in August, you have three weekends left. That is enough to write a procedure, decide who picks up the phone and publish a security contact address. It is not enough to reconstruct the subcontracting chain of firmware shipped five years ago, or to retrofit an update mechanism onto an installed base that never had one. Do the part that fits on one page now; the rest has until 2027.

And if you are still unsure which side of the test you fall on, that is already an answer: nobody should discover their regulatory status on the day the flaw gets exploited. It takes one conversation to settle, and our Cyber Resilience Act compliance pages set out the framework.