# 24 Hours: Can Europe Keep the Secrets It Now Collects?

6 NOVEMBER 2026

**Andrea Bianchi**



Photo by Christian Lue on Unsplash

Since 11 September 2026, any manufacturer of a connected product sold in the European Union has had 24 hours to warn the authorities once it learns that a flaw in that product is being actively exploited. That is the core of Article 14 of the Cyber Resilience Act, which applies more than a year before most of the regulation does, in December 2027. A fuller notification follows within 72 hours, and a final report within 14 days of a fix becoming available.

The infrastructure behind it arrived only just in time. ENISA published the list of coordinating national CSIRTs on 4 September and the platform's address and user documentation on 10 September, a day before the obligations applied. The Single Reporting Platform went live on the deadline itself.

The aim of this piece is not to question whether Europe should report vulnerabilities quickly; I think it should. It is to ask whether the EU has been as careful about protecting that information as it has been about collecting it. In my view, it has not. By routing every early warning about Europe's

most dangerous live vulnerabilities through a single system, the EU has built something adversaries will want badly, and the safeguards around it deserve far more scrutiny than they have had.

## Why speed is worth it

The case for 24 hours is simple: attackers do not wait. In the first half of 2026, CrowdStrike found that 88% of the exploitation it observed against vulnerabilities with public proof of concept code began within 48 hours of that code's release. A reporting regime built around the patch cycle would always arrive after the damage.

Article 14 also fixes a structural blind spot. Before September, a manufacturer whose router was being exploited in five member states had no EU wide legal duty to report the flaw. Now one notification reaches the coordinating CSIRT and ENISA at the same time, giving Brussels a continental picture of what is being exploited and where.

There is a foreign policy dimension here that is easy to miss. In April 2025, the CVE programme, funded by the US government and the backbone of global vulnerability tracking, came within hours of losing its contract before CISA stepped in at the last minute. A month later, ENISA launched the European Vulnerability Database, built to serve the NIS2 Directive and the Cyber Resilience Act but suddenly very timely. Article 14 follows the same logic: Europe should not depend on a single foreign agency's budget to know which of its own products are under attack. Vulnerability intelligence has become a question of strategic autonomy, not only of compliance.

## The price of a single channel

None of this was uncontroversial. In October 2023, around fifty security experts signed an open letter warning that the draft CRA would give dozens of government bodies real time access to unpatched flaws, creating a target for criminals and a temptation for intelligence services. A coalition of more than thirty industry and technology associations, including DIGITALEUROPE and BSA, made a similar point: concentrating such data is a risk in itself.

The final text listened in part. Notifications go to a national CSIRT, which may in exceptional cases delay sharing them with other member states, and in rarer cases flagged by the manufacturer, ENISA at first receives only basic details. Delegated Regulation (EU) 2026/881 sets out when a CSIRT may delay, including where the platform itself has been compromised. But delay is the exception: by default, a notification travels to the CSIRT of every member state where the manufacturer says the product is available. And the delegated act's rule for a compromised platform tells us how Brussels would contain a leak, not how it would prevent one. Prevention rests on a duty in Article 16 of the CRA for ENISA to take appropriate and proportionate

security measures, and outsiders have little way to judge how well that duty is met.

The same question is now being fought out on a larger scale. In its Digital Omnibus of November 2025, the Commission proposed one EU entry point for incident reports under NIS2, GDPR, DORA and other laws, run by ENISA and building on the CRA platform. Member states have raised concerns that it could become a single point of failure, and on 21 October the Council replaced it in its negotiating position with a national entry point in each member state. The European Parliament has yet to adopt its own position. I think the Council's instinct is sound. One ENISA system holding exploited vulnerability reports alongside personal data breaches and financial sector incidents would be among the most valuable intelligence targets in Europe, for state actors as much as for criminals.

## Keep the clock, guard the channel

The answer is not to slow the clock down. It is to treat what the clock produces as sensitive intelligence. Three changes would help.

First, hold the platform to the letter of the law. The CRA asks only for the general nature of an exploit at the 72 hour stage, and requires strict need to know handling where no fix exists. ENISA's templates and practice should keep it that way, so that detail precise enough to reproduce an exploit never sits on the platform before a fix is available.

Second, make the platform's security verifiable. ENISA should commission regular independent penetration testing and publish an annual summary, including how often CSIRTs invoked the delay grounds under the delegated act, and why.

Third, keep the stores separate. Whatever form the Omnibus finally takes, vulnerability data and personal data breach reports should not sit in one database simply because they could share a front door.

A day is a reasonable thing to ask of manufacturers. The real test is whether the EU will hold itself to an equally demanding standard for protecting what it now collects. As the Omnibus negotiations continue, that is a question Brussels should answer before the first breach, not after it.