



EU Unpacked

www.euawareness.org

#KnowledgeAgainstManipulation

EU Unpacked # 14

Dr George Robakidze

The Cyber Resilience Act

An In-Depth Analysis for Non-Experts

The opinions and conclusions expressed are those of the author and do not necessarily reflect the views of the EU Awareness Centre.

September 2026

© [2026] EU Awareness Centre. All rights reserved.

I. Closing the Gap: Why the EU Legislated on Product Cybersecurity

The gap the Regulation was built to close

For most of the last two decades, European cybersecurity law addressed organisations rather than objects. The [NIS2 Directive](#) tells hospitals, energy operators and digital infrastructure providers how to manage risk. [The 2019 Cybersecurity Act](#) built a certification scheme that companies could use if they chose to. Neither instrument, however, told the manufacturer of a smart thermostat, a home router or an industrial sensor that its product had to be secure before it reached a shop shelf. That gap sat uneasily beside a fact regulators had long known and citizens experienced directly: cheap, badly maintained connected devices were becoming the easiest way into European homes, offices and, on occasion, critical infrastructure.

The **Cyber Resilience Act (CRA)** — formally [Regulation \(EU\) 2024/2847](#) of 23 October 2024 — is the European Union's answer to that gap. It is a horizontal regulation, meaning it does not target one sector but reaches across the entire market for **products with digital elements**: any piece of hardware or software, sold commercially in the Union, whose function depends on a data connection to a device or a network. The recitals accompanying the Regulation are candid about the scale of the problem it addresses. They describe "a low level of cybersecurity" reflected in "widespread vulnerabilities" and "insufficient and inconsistent provision of security updates," compounded by a second failure: ordinary buyers had no reliable way of judging which products were safe to use. Cyberattacks, the recitals note, are no longer merely a commercial nuisance; they are treated as a matter of public interest with consequences for democracy, consumer safety and health.

What makes the CRA structurally significant is not the ambition of that diagnosis — most cybersecurity strategies open with similar language — but the choice of remedy. Rather than asking Member States to regulate products individually, or leaving the market to self-correct through liability suits after the damage is done, the Union chose to impose uniform, ex-ante design obligations directly on manufacturers, enforced through the same CE-marking architecture used for machinery, toys and medical devices. A washing machine cannot legally be sold in the EU unless it meets safety standards decided in Brussels and Strasbourg; from 2027, in principle, neither can an insecure smart washing machine.

A single rulebook for a fragmented threat

The case for horizontal legislation rested on a simple observation: national approaches to product cybersecurity, where they existed at all, differed enough to fragment the single market without meaningfully reducing risk. A manufacturer selling into twenty-seven jurisdictions faced twenty-seven possible sets of informal expectations and no reliable European baseline against which to design a product once. The Regulation's own explanatory memorandum frames this as "a legislative patchwork within the internal market," and the phrase captures the underlying political logic well. This is less a story about cybersecurity idealism than about the EU's oldest instinct: fragmentation is treated as a cost in itself, independent of the underlying risk it is meant to manage.

That instinct explains why the CRA was built on Article 114 of the Treaty on the Functioning of the European Union — the internal market clause — rather than on a security-specific legal basis. The choice is not merely technical. It signals that Brussels regards product cybersecurity primarily as a market-functioning question, subject to qualified-majority voting and full co-decision with the European Parliament, rather than as an area where Member States might insist on unanimity or a narrower security competence. The same legal-basis question, as later sections show, has become considerably more contentious in a related but separate piece of legislation now moving through the institutions.

The result is a Regulation that reads, in places, like an industrial-safety statute transplanted into the digital domain. That transplant is the CRA's defining feature: it treats insecure code as a defective product, not as an externality to be managed after the fact.

II. Scope and Definitions: What Counts as a Product With Digital Elements

Drawing the boundary of the market

The Regulation's scope is deliberately broad. A **product with digital elements** covers standalone software such as mobile applications, hardware with embedded software such as routers and connected cameras, standalone hardware including integrated circuits, and any combination of the two intended to work together. Article 2 draws the line at commercial availability: if a product is placed on the EU market — supplied for distribution or use in the course of a commercial activity, whether paid or free — it falls within scope, provided its intended or reasonably foreseeable use includes a logical or physical connection to a device or network.

That single criterion excludes more than it might first appear to. A website that does not itself perform the function of a connected product falls outside the Regulation. Standalone software-as-a-service offerings, cloud platforms built and operated independently of any physical product, generally sit outside the CRA's scope too, unless they qualify as **remote data processing** — a service designed by, or under the responsibility of, the manufacturer, without which the product could not perform one of its core functions. A smart doorbell whose motion-detection feature depends on a manufacturer-run cloud service brings that cloud service inside the Regulation's reach; a general-purpose video storage app the user separately chooses to connect to the same doorbell does not.

The practical effect is to make the manufacturer's supply-chain architecture, not merely its physical product, a matter of legal consequence. A firm that designs its cloud dependency out of the product removes an entire category of obligation; one that designs it in inherits responsibility for a service that may run for years after the physical device has shipped.

Who bears the obligations?

The Regulation distributes responsibility across what it calls **economic operators**: manufacturers, authorised representatives, importers and distributors. Manufacturers carry the overwhelming share of substantive obligations, because they are best placed to design security in from the start; importers and distributors act principally as gatekeepers, obliged to verify that a product already carries the required documentation and CE marking before they place it on the market or make it available. This graduated model — heaviest duties on the party closest to design decisions, lighter but real duties further down the chain — mirrors the structure the EU already uses for physical product safety, and it is precisely that continuity with familiar law that gives the CRA its practical bite: national market-surveillance authorities do not need to invent new enforcement doctrine, only extend an existing one to a new class of product.

One definitional innovation deserves particular attention because it shaped the entire legislative debate: the **open-source software steward**, a category created specifically to avoid treating every open-source foundation as a commercial manufacturer. Its introduction, discussed in full in Section VI, was the price of reconciling a product-liability logic built for factories with a development model built on volunteers.

Taken together, the scope provisions do something more consequential than list exclusions. They establish that European law now treats "placing a connected product on the market" as an act carrying default legal consequences, reversing a long-standing assumption that digital products were free to enter commerce unless a specific rule said otherwise.

III. The Manufacturer's Core Obligations: Security by Design and by Default

Risk assessment as the organising discipline

Article 13 is the operational heart of the Regulation. It requires manufacturers to ensure, before placing a product on the market, that it has been designed, developed and produced in accordance with the essential cybersecurity requirements set out in Annex I. Compliance is not judged against a fixed checklist alone; it begins with a mandatory **cybersecurity risk assessment** that must analyse the product's intended purpose, its reasonably foreseeable use, its operational environment, and the length of time it is expected to remain in service. That assessment then has to be documented, kept in the product's technical file, and revisited throughout the support period rather than filed away once at launch.

This is a meaningfully different regulatory philosophy from a simple pass/fail certification test. It asks manufacturers to reason continuously about threat exposure rather than to demonstrate a one-off snapshot of compliance, and it makes the quality of that reasoning — not merely its existence — into something a market-surveillance authority can later scrutinise. A manufacturer that conducts a superficial risk assessment and one that conducts a rigorous one may both tick the same procedural box, but only one will hold up when a regulator asks to see the file after an incident.

Consider a mid-sized European manufacturer of smart building-access controllers. Under the CRA, it cannot simply ship a device with a default password and a promise of "updates as needed." It must assess how the device will realistically be deployed — likely on a building network with mixed security maturity, likely left in place for a decade — and design accordingly, then commit that reasoning to a document a regulator can later inspect.

Due diligence through the supply chain

Modern software is assembled, not written from scratch, and the Regulation takes that reality seriously. Article 13(5) requires manufacturers to exercise due diligence when integrating third-party components — including free and open-source components not themselves placed on the market — so that those components do not compromise the security of the finished product. Where a manufacturer discovers a vulnerability in an integrated component, it must notify whoever develops or maintains that component and, where it has produced a fix, share the relevant code or documentation with them.

This due-diligence duty is where the CRA quietly rewires incentives across the entire software supply chain. A manufacturer that has never previously audited its dependencies now has a legal reason to do so, and a legal reason to feed fixes back upstream rather than patch only its own downstream copy. The mechanism is ambitious, because it tries to fix collective-action failures that have plagued open-source security for a decade — the fact that everyone relies on shared components and almost no one funds their maintenance. It is also fragile, because the obligation to exercise due diligence says nothing about who pays for the audits, the disclosures or the fixes it generates.

The manufacturer's obligations, read together, describe a shift from cybersecurity as an optional feature sold at a premium to cybersecurity as a precondition of market access — a shift the Regulation's own explanatory memorandum calls, without exaggeration, foundational.

IV. Vulnerability Handling, Support Periods and the Reporting Clock

How long must a product remain safe to use

One of the CRA's most consequential — and most argued-over — provisions is its requirement that manufacturers determine a **support period**: the length of time during which they commit to handling vulnerabilities in a product. Article 13(8) requires that this period reflect realistic user expectations and the

product's likely lifespan, and Commission guidance issued alongside the Regulation has since clarified that the commonly cited five-year floor is exactly that — a floor, not a default. A product reasonably expected to remain in use for longer, such as an industrial sensor or a piece of building infrastructure, is expected to carry a correspondingly longer support commitment. Once fixed, the end date of that period has to be disclosed to buyers at the point of purchase, and the underlying technical documentation must remain available for at least ten years — or the length of the support period if that is longer, whichever gives users the greater protection.

This converts what has historically been a marketing choice — how long a vendor bothers to keep supporting a device — into a disclosed, legally anchored commitment. It is difficult to overstate how large a change that is for an industry accustomed to abandoning products the moment a newer model ships.

The reporting clock: 24 hours, 72 hours, and after

From 11 September 2026, manufacturers must notify any actively exploited vulnerability and any severe incident affecting a product's security to the national **Computer Security Incident Response Team (CSIRT)** and to [ENISA](#), the EU cybersecurity agency, simultaneously and through a single reporting platform. The clock is unforgiving by design: an early warning within twenty-four hours of becoming aware of the problem, a fuller notification within seventy-two hours, and a final report within fourteen days of a fix becoming available for an exploited vulnerability, or within one month of the seventy-two-hour notification for a severe incident.

These deadlines were deliberately aligned with the equivalent timelines under NIS2, a decision several industry bodies specifically welcomed, since a single incident affecting both a product manufacturer and the operator using it would otherwise trigger two different reporting clocks running on two different schedules. Alignment here is not a cosmetic courtesy; it is what makes simultaneous compliance administratively survivable for firms that sit inside both regimes.

A short transitional exception softens the edges for the smallest firms: micro and small enterprises cannot be fined specifically for missing the twenty-four hour deadline, though the underlying obligation to report still applies. That concession, small as it looks on the page, was one of the more hard-fought details in the legislative process, and it illustrates a recurring pattern in this Regulation — a rule written for the market as a whole, softened at the margins for the actors least equipped to absorb it.

The overall architecture treats a vulnerability disclosed and fixed within days as a manageable cost of doing business, and a vulnerability concealed or left unaddressed as a regulatory failure in its own right — a distinction with real consequences once the enforcement provisions, discussed in Section VII, are applied.

V. Conformity Assessment and Market Access: Self-Certification, Notified Bodies and CE Marking

A tiered system of scrutiny

Not every connected product carries the same risk, and the CRA does not treat them as though it did. The Regulation sorts products into a default category and two elevated categories — broadly, "important" products with digital elements (Class I and Class II) and a narrower band of "critical" products — based on the function they perform and the consequences of their failure. Password managers, firewalls, and operating systems for critical infrastructure sit in the more heavily scrutinised tiers; a smart kettle, in the ordinary run of cases, does not.

For the default category, manufacturers may self-assess their own conformity with Annex I's essential requirements. For the elevated tiers, the Regulation requires an external check — either through harmonised standards applied in full or, in the absence of adequate standards, through third-party assessment by a

notified body, a conformity-assessment organisation designated by a Member State and recognised across the Union. Once a product has cleared the relevant assessment, the manufacturer draws up an EU declaration of conformity and affixes the CE marking — the same visual signal European consumers already associate with physical product safety, now extended to claim something about a product's cybersecurity as well.

The standards problem hiding behind the marking

The CE marking's credibility depends entirely on the existence of harmonised technical standards that give Annex I's broad principles concrete, testable content — and those standards have lagged behind the legislative timetable. This is not an incidental delay. It is the reason the Digital Omnibus on Artificial Intelligence proposal, moving through the institutions alongside the CRA, explicitly links the entry into application of comparable high-risk obligations under the AI Act to the actual availability of supporting standards rather than to a fixed calendar date. The Commission's own guidance acknowledges the same tension for the CRA: manufacturers are expected to demonstrate conformity with requirements for which the detailed technical yardstick is, in places, still being written.

The consequence is a period — likely to run through 2026 and into 2027 — in which conformity assessment will rest more heavily on manufacturers' own technical judgement, informed by draft standards and Commission FAQs, than the architects of the Regulation originally intended. That is not a fatal flaw. It is, however, a reminder that a CE mark is only as reliable as the standards infrastructure behind it, and that infrastructure takes longer to build than a legislative text takes to adopt.

The tiered assessment model is the clearest expression of a principle running through the whole Regulation: proportionality by consequence, not by sentiment. A product's risk tier, not the sympathy its manufacturer commands, decides how much scrutiny it receives.

VI. The Open-Source Compromise: Stewards, Not Manufacturers

A backlash that changed the text

No part of the CRA generated more sustained controversy before adoption than its treatment of free and open-source software. An early 2022 draft exempted open-source software developed "outside the course of a commercial activity," language that alarmed the open-source community for reasons that were more than semantic. **OpenForum Europe**, the **Open Source Initiative** and a coalition of foundations warned that the exemption's wording did not match how open-source projects actually function — through sponsorship, bug bounties, corporate contributions and informal support arrangements that blur any clean line between "commercial" and "non-commercial" activity. Simon Phipps of the OSI went further, suggesting that developers facing uncertain liability might simply geo-block European users rather than risk exposure — a scenario in which a regulation meant to make Europeans safer would instead cut them off from globally maintained security fixes.

The final text responds to that backlash with a purpose-built category: the **open-source software steward** — typically a foundation or similar legal person that provides sustained, systematic support for the development of open-source products intended for commercial use, without itself manufacturing or selling them. The Linux Foundation and the Apache Software Foundation are the paradigm cases the legislative debate had in mind. Stewards face a lighter, bespoke regime: a documented cybersecurity policy, cooperation with market-surveillance authorities, and vulnerability reporting only to the extent they are actually involved in a product's development. Critically, Article 64(10) exempts them from the administrative fines that apply to manufacturers — an explicit acknowledgement that treating a nonprofit foundation the same as a commercial vendor would have been disproportionate, and quite possibly self-defeating.

An unresolved tension, not a closed question

The compromise is more workable than the original draft, but it has not silenced the underlying debate. Academic assessments published after adoption — including analysis in the *Journal of Intellectual Property, Information Technology and E-Commerce Law* — conclude that the CRA "does achieve a balance between regulation for open-source software and advancing cybersecurity, but only through a narrowly applicable and, at times, complex legislative approach." Complexity is doing real work in that sentence: a steward's obligations depend on the specific nature of its involvement in a given product, a fact-sensitive test that smaller foundations, lacking in-house legal counsel, may struggle to apply with confidence.

A separate strand of scholarship reads the open-source provisions through the lens of **digital sovereignty**, arguing that the EU's instinct to assert more control over software supply chains sits uneasily with a development model built on openness and voluntary global collaboration. That critique does not claim the CRA is hostile to open source — the Regulation's recitals go out of their way to praise its contribution to innovation — but it does suggest that the Union's regulatory reflexes and open source's operating logic will keep rubbing against each other as implementation proceeds, particularly wherever due-diligence obligations reach into unpaid, informal contributions.

The open-source chapter of this story illustrates something the CRA cannot resolve by drafting alone: a law built on the logic of manufacturers and products will always sit at an angle to a development culture built on contributors and commons, however carefully the definitions are drawn.

VII. Enforcement and the Price of Non-Compliance

Fines calibrated to deter, not merely to punish

The CRA leaves the detailed design of penalties to Member States but sets a ceiling meant to concentrate minds at board level: administrative fines of up to €15 million or 2.5% of a company's total worldwide annual turnover for the preceding financial year, whichever is higher. That formula — familiar from the GDPR — is a deliberate signal that Brussels intends product cybersecurity to be treated with the same seriousness as data protection, not as a lesser technical annexe to it.

Financial penalties sit alongside more immediately disruptive tools: national market-surveillance authorities can require a product's withdrawal from sale, demand a recall, or restrict its availability pending corrective action. For a manufacturer, a recall order is frequently more damaging than the fine that might accompany it — reputational harm and lost shelf space compound faster than a one-off financial penalty, and that asymmetry is precisely what gives market-surveillance powers their practical force.

Who actually polices this

Enforcement authority sits primarily with national market-surveillance bodies, which can demand technical documentation, order testing, and act on complaints or on their own initiative. ENISA and the CSIRT network play a coordinating role, principally through the vulnerability and incident reporting channel established under Article 14. This is a genuinely federated enforcement model: the rulebook is set in Brussels, but the inspectors, for the most part, work for national governments. That division of labour has served EU product-safety law reasonably well for physical goods; whether twenty-seven national authorities, of uneven size and technical capacity, can apply it with comparable rigour to fast-moving software products is one of the Regulation's genuine open questions, and one that will only be answered once the first serious enforcement cases begin.

The enforcement architecture, taken as a whole, tells manufacturers something simple and consequential: cybersecurity failures are no longer purely reputational or commercial risks to be managed internally. They

are now regulatory risks, carried on the same institutional rails — and increasingly the same order of financial magnitude — as a breach of the Union's data protection law.

VIII. A Long, Staggered Road: The CRA's Phased Timeline

Four dates that matter

The Regulation entered into force on 10 December 2024, but "in force" and "in application" are different things in EU law, and the gap between them here is unusually long and deliberately staggered. Four dates define the transition. Governance provisions on the notification of conformity- assessment bodies apply from 11 June 2026, so that the infrastructure needed to certify higher-risk products exists before manufacturers are required to use it. Reporting obligations under Article 14 — the twenty-four and seventy-two hour vulnerability and incident notification duties described in Section IV — apply from 11 September 2026, a date that, as this analysis is being written, has only just passed. The great majority of substantive obligations, including the essential cybersecurity requirements and the conformity-assessment procedures described in Sections III and V, apply from 11 December 2027. A narrower category of legacy products already on the market before that date remains outside the Regulation's reach unless it later undergoes a substantial modification, though reporting duties apply to them regardless.

Staggering the timeline this way was a considered choice, not an accident of drafting. It gives standards bodies, notified bodies and national authorities time to build the infrastructure the Regulation assumes exists, while bringing the single most urgent obligation — telling someone when a product is actively being exploited — into force years ahead of the rest.

A deadline under quiet strain

Even so, industry analysis published around the September 2026 reporting deadline describes a compliance landscape still under construction. Harmonised standards for many of Annex I's essential requirements remain incomplete; the Commission has responded with interim practical guidance rather than finished technical norms, most recently a guidance document published on 27 July 2026, weeks before the reporting obligation took effect. That sequencing — guidance arriving just ahead of a legal deadline, rather than comfortably in advance of it — has become a recognisable feature of large EU digital-regulation projects, and the CRA is not the exception to that pattern some had hoped it would be.

None of this means the timeline will slip in a formal sense; the dates written into the Regulation are binding law, not aspirations. It does mean that the practical experience of compliance in the second half of 2026 will look less like following a finished rulebook and more like applying general principles under continuing uncertainty about their precise technical content — a condition manufacturers of physical goods rarely have to tolerate, and one that partly explains the compliance-burden debate examined next.

IX. Simplification Fatigue: SMEs, Competitiveness and the Politics of Compliance Burden

What the data on small manufacturers actually shows

A 2026 ENISA survey of small and medium-sized manufacturers, conducted as part of the Commission's SME cybersecurity strategy, offers an unusually direct look at how the CRA lands on the businesses least equipped to absorb it. Two-thirds of respondents said they were already familiar with the Regulation, but that familiarity did not translate into confidence: technical documentation and conformity assessment were named as the weakest areas of understanding, and financial support was the single most requested form of assistance. More than a third of microcompanies reported having no incident- response plan at all — precisely the capability the CRA's reporting clock now assumes they possess.

That gap between legal obligation and organisational capacity is not a marginal detail; it is the central practical risk to the Regulation's credibility. A rule that assumes a level of security maturity most micro-enterprises do not yet have will either force genuine investment in that maturity — the outcome its architects intend — or push the smallest, least resourced manufacturers toward under-reporting, informal non-compliance, or exit from the EU market altogether. Which of those outcomes predominates will not be clear for several years.

A regulation caught in a larger political current

The CRA's compliance burden has also become a proxy in a much larger fight over the direction of EU regulatory policy. Business associations such as **Eurochambres**, representing chambers of commerce across the Union, have welcomed the Regulation's objectives while pressing for stronger safeguards against "gold-plating" by Member States and clearer alignment between the CRA, NIS2 and the Cybersecurity Act, so that a single company does not face three overlapping reporting duties for a single incident. Commentary aimed at a general business readership has been considerably less measured, framing the CRA as one more entry in a lengthening list of rules that, in combination, risk pricing small European technology firms out of their own market — the same argument, in essence, that animates the Commission's own Digital Omnibus simplification package.

That the Commission is simultaneously defending the CRA's substance and promoting a horizontal simplification agenda elsewhere in the digital rulebook is not a contradiction so much as a reflection of where European politics has moved since the Regulation was first proposed in 2022. A measure conceived at the height of confidence in the EU's capacity to set global standards through regulation is now being implemented in a political climate considerably more anxious about competitiveness — a shift this publication has traced across several of the acts examined in this series, from the Digital Services Act to the Critical Raw Materials Act.

The compliance-burden debate will not be settled by this analysis, and it should not be. What can be said with confidence is that the CRA's ultimate reputation — ambitious safeguard or needless burden — will be decided less by its text than by whether the support measures promised to smaller manufacturers, financial and technical alike, actually materialise before the December 2027 deadline arrives.

X. Where the CRA Sits in the EU's Cybersecurity Architecture

Product rules alongside organisational rules

The CRA does not operate alone. It is the product-level complement to NIS2's organisational requirements and the Cybersecurity Act's voluntary certification framework — together, the three instruments are intended to cover, respectively, what a product must be, how an organisation using it must behave, and how either can seek independent certification of its security claims. The reporting-deadline alignment discussed in Section IV is the clearest evidence that the Commission designed these instruments to interlock rather than to compete, even though they were adopted at different times through different legislative processes.

That architecture is currently being extended again. A proposed overhaul of the Cybersecurity Act — informally labelled **CSA2** — would give the Commission new powers to designate "high-risk third countries" and exclude named suppliers from sensitive ICT supply chains, extending the logic of the 2019 5G Toolbox from a voluntary recommendation into a binding mechanism. Legal commentary in *Euractiv* has flagged a serious institutional question about that proposal: built on the same internal-market legal basis as the CRA itself, it may in substance be pursuing an objective — national security — that Article 4(2) of the Treaty on European Union reserves to Member States alone. Should that argument prevail before the Court of Justice, it would not touch the CRA directly, but it illustrates how contested the outer edge of "internal market"

competence has become in the cybersecurity domain generally, and it is worth noting because CSA2 is frequently — and inaccurately — conflated with the CRA in general commentary. They are related instruments pursuing different objectives through different mechanisms, and readers should resist collapsing them into a single "EU cybersecurity package."

The AI Act's shadow

A further point of contact runs through the **Artificial Intelligence Act**, examined earlier in this series. Where a product with digital elements also qualifies as a high-risk AI system, its cybersecurity risk assessment under the CRA may be integrated into the risk assessment already required by the AI Act, rather than duplicated. That interplay matters practically for any firm building connected products with embedded AI functionality — increasingly the industry norm rather than the exception — and it is one of the more genuinely simplifying features of a regulatory landscape more often criticised for adding layers than removing them.

Read together, these overlapping instruments describe an EU cybersecurity architecture still very much under construction, its pieces added at different moments in response to different anxieties, now being stitched together — not always smoothly — into something its architects hope will eventually function as a coherent whole rather than a collection of adjacent rules.

XI. Relevance for EU-Aspiring Countries

A rulebook candidate countries will eventually inherit

For the nine countries currently negotiating or awaiting EU membership — Ukraine, Moldova, Georgia, and the six Western Balkan states of Albania, Bosnia and Herzegovina, Kosovo, Montenegro, North Macedonia and Serbia — the Cyber Resilience Act is not yet a direct legal obligation, but it is a preview of the *acquis* each of them will eventually be expected to transpose. The European Union Institute for Security Studies noted in a 2026 policy brief that digital and cybersecurity alignment across the enlargement process remains "largely embedded in a gradual, compliance-driven accession process," with regulatory transposition often running ahead of the operational capacity needed to make the rules meaningful in practice. That warning applies to the CRA with particular force, since the Regulation assumes an ecosystem of notified bodies, national market-surveillance authorities and CSIRTs that took the EU-27 itself years to build, and that no current candidate country yet possesses in comparable form.

Legal starting points across the group vary considerably, and the variation matters for how each will eventually approach transposition. EU4Digital assessments of the Eastern Partnership found Georgia to be the only country in that group to have completed a third iteration of its national cybersecurity strategy, ahead of Moldova and Ukraine, which were still updating theirs, and well ahead of Armenia, Azerbaijan and Belarus, which had not yet adopted one at all. The Western Balkan candidates present a different picture again: several have transposed elements of the original NIS Directive over the past decade under the Energy Community and stabilisation processes, giving them a longer, if uneven, history of engagement with EU cybersecurity law than most Eastern Partnership states. None of this variation amounts to CRA-readiness anywhere in the group. Every candidate's existing framework, whatever its maturity, remains built around organisational or critical-infrastructure obligations rather than the EU's newer, more granular approach to product-level conformity assessment, support periods and CE marking for software — concepts that did not exist in European law when most of these national frameworks were drafted.

Capacity before compliance

The EU's own technical assistance instruments implicitly recognise this gap, and they are organised regionally rather than around any single candidate. The **Western Balkans Cyber Capacity Centre** supports the six Balkan candidates on operational cooperation and joint exercises; the **Eastern Partnership Cybersecurity Initiative**,

delivered through the e-Governance Academy in Georgia, Azerbaijan and Moldova alongside GIZ's parallel work in Ukraine and Armenia, has deliberately emphasised digital literacy and civil-society engagement — particularly outside capital cities — alongside the more conventional work of legal alignment. Both reflect the same underlying judgement: legislative transposition without a corresponding increase in institutional capacity produces compliance on paper without resilience in practice. The EUISS brief drew the identical conclusion at the regional level, arguing that enlargement partners across both the Western Balkans and the Eastern Neighbourhood should be treated as "frontline actors" with operational capabilities worth integrating into EU structures before accession, rather than as passive recipients of an *acquis* to be copied once membership talks conclude.

For any country in this position, the CRA's most immediate lesson is probably not its specific technical requirements — those will matter only once accession negotiations reach the relevant chapter — but the institutional model behind them: a functioning system of conformity assessment, credible national authorities empowered to inspect and sanction, and a genuine reporting culture between manufacturers and government. Building that institutional muscle now, ahead of any formal transposition deadline, would do more for accession readiness across the group than a premature attempt to mirror the CRA's text in national law. **Georgia** illustrates the point sharply: a state contending with a sustained pattern of Russian-linked cyber activity against its own institutions has good reason to build this capacity regardless of the accession timetable, not merely because Brussels will eventually ask it to — and the same logic, in different geopolitical registers, applies to Ukraine's wartime digital infrastructure and to the Western Balkan states' exposure to hybrid pressure from outside the region.

Conclusion. A Structural Shift, Arriving Under Strain

What the CRA actually changes

The Cyber Resilience Act does something European law had not previously attempted at this scale: it makes the absence of basic cybersecurity, in a commercial product sold into the EU market, a legal defect rather than a business risk the market was left to price in on its own. That is a genuine structural shift, comparable in kind — if not yet in cultural weight — to the moment product-safety law first made an unsafe toy or a faulty appliance illegal to sell rather than merely unwise to buy. Manufacturers, importers and distributors now operate under a default assumption that would have been unfamiliar a decade ago: cybersecurity is not a feature to be marketed at a premium, but a condition of lawful market access.

The Regulation's architecture is coherent and, in its broad strokes, defensible. Risk-based obligations scale with the harm a product's failure could cause. Reporting deadlines are aligned with neighbouring legislation rather than invented in isolation. A genuinely difficult political problem — how to regulate commercial products without destroying an open-source ecosystem built on entirely different premises — was met with a purpose-built compromise rather than an afterthought.

None of that guarantees the CRA will work as intended. Standards remain incomplete as binding deadlines approach. Small manufacturers report confidence gaps in exactly the areas the Regulation treats as foundational. And the Regulation now takes effect inside a political environment that has grown considerably more sceptical of the regulatory instinct that produced it, with the Commission's own simplification agenda circling adjacent parts of the digital rulebook even as the CRA's first hard deadlines land. Whether the Cyber Resilience Act is remembered as the moment European cybersecurity law matured, or as a well-designed rule that arrived a political cycle too late for the appetite to enforce it fully, will depend less on the text analysed in this piece than on choices European institutions have not yet made — about funding, about standards, and about how much regulatory ambition the Union's current mood can still sustain.

¹ **Regulation (EU) 2024/2847 (Cyber Resilience Act)**

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, OJ L, 2024/2847, 20.11.2024. Available at [EUR-Lex](#).

² **European Commission, 'The Cyber Resilience Act — Summary of the legislative text'**

Shaping Europe's Digital Future, digital-strategy.ec.europa.eu, updated 3 December 2025.

³ **Pinsent Masons, 'Cyber Resilience Act marks major shift in EU cybersecurity compliance'**

Out-Law analysis, 25 August 2026, summarising Commission guidance published 27 July 2026 on manufacturer obligations, support periods and reporting deadlines.

⁴ **ENISA, 'SME CRA Survey Report'**

European Union Agency for Cybersecurity, June 2026, summarising SME awareness of and preparedness for Cyber Resilience Act obligations.

⁵ **Eurochambres, 'Position on the Cybersecurity Package'**

Association of European Chambers of Commerce and Industry, June 2026.

⁶ **OpenForum Europe et al., submission on the 2022 Cyber Resilience Act draft**

Joint submission by OpenForum Europe, the Eclipse Foundation, the Open Source Initiative and others to the European Commission, January 2023.

⁷ **JIPITEC, 'The Cyber Resilience Act and Open-Source Software: A Fine Balancing Act'**

Journal of Intellectual Property, Information Technology and E-Commerce Law, Vol. 16 No. 1 (2025).

⁸ **Euractiv, 'Brussels' proposed cybersecurity overhaul risks constitutional red flags'**

Euractiv, 4 March 2026, on the proposed revision of the Cybersecurity Act ('CSA2').

⁹ **European Union Institute for Security Studies, 'Shared threats, shared resilience'**

EUISS Brief, 26 May 2026, on integrating enlargement partners into EU digital and cyber security frameworks.

¹⁰ **EU4Digital / e-Governance Academy, Eastern Partnership Cybersecurity Initiative**

e-Governance Academy project documentation and EU4Digital, 'Cybersecurity guidelines for the Eastern Partner countries', on the state of cybersecurity legislation in Georgia and other Eastern Partnership countries.

Dr George Robakidze



George Robakidze is a diplomat and policy expert specialising in international politics, security, and European integration. He has held senior roles in Georgia's Ministry of Foreign Affairs, focusing on political affairs, Euro-Atlantic integration, and regional security.

He holds a PhD in Political Science. His research focuses on radical and populist movements in Eastern Europe, alongside broader political and international issues.