Open:UK®

Cyber Regulation Report 2026

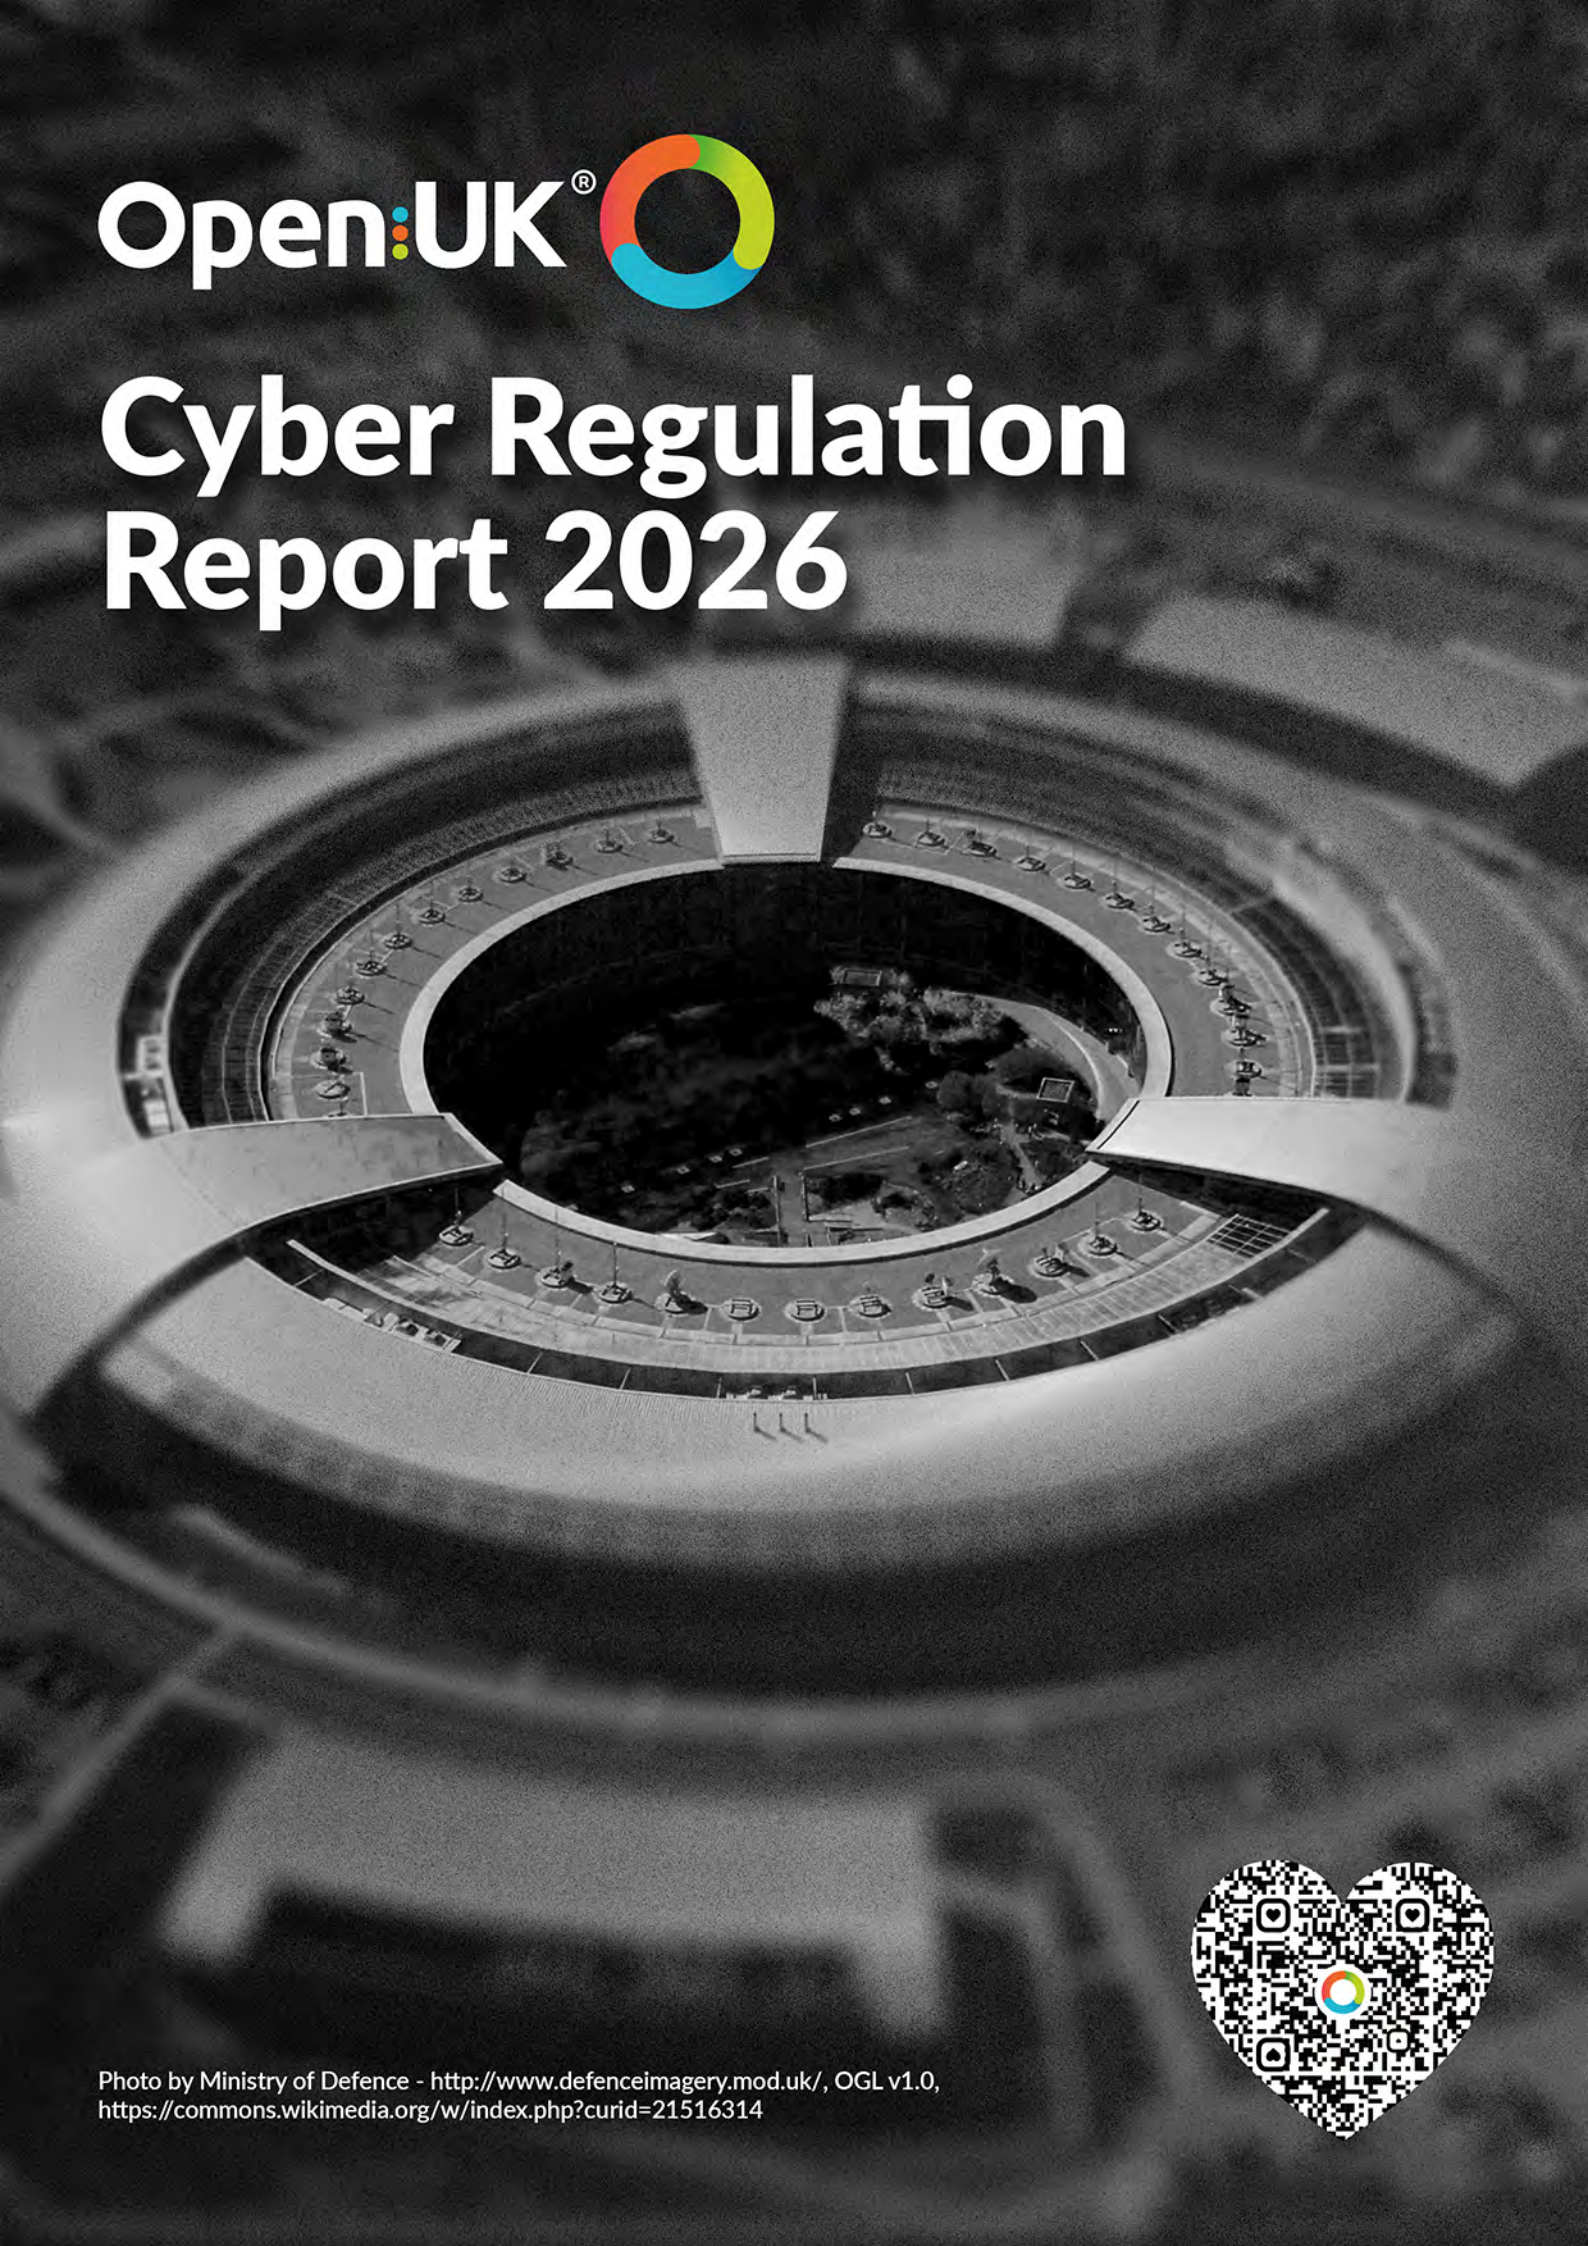

Index

1. Introducing the Cyber Report

1.1 Executive Summary: TL;DR

1.1.1 Overview

Open source software underpins the modern digital economy, present in almost every codebase, yet its security has remained a largely unregulated space. This Report examines how cyber regulation worldwide is, or sometimes is not, adapting to that reality, and how the arrival of frontier AI is reshaping the risk landscape faster than the law can follow. It is organised into five parts.

1.1.2 The regulatory landscape and the open source blind spot

The Report opens by mapping how cyber regulation divides into two families: **organisation-centric regimes** that impose duties on operators of critical infrastructure (NIS2, Singapore's Cybersecurity Act, India's CII regime, and the UK's forthcoming Bill), and **product-centric regime**s that attach obligations to software itself. Of the twenty-three jurisdictions surveyed, only the EU and the US substantively address open source at all in regulation. Elsewhere, pressure reaches maintainers indirectly through supply-chain obligations that flow upstream without ever naming open source creating a regulatory perimeter built around the ecosystem rather than with it, despite community exclusions due to commercialisation provisions.

Most recently, President Trump has signed a National Security Presidential Memorandum that aims to put cutting edge AI tools into the hands of the US military. The Trump administration is establishing another framework that would "accelerate AI adoption" across a network of federal defense agencies and "adapt the best commercial and open source technologies for mission use."

The Memo does not say that national security AI must be proprietary, nor does it suggest restricting open source models. It explicitly directs agencies to leverage open source technologies when they are the best tool for the mission. The Executive Order published shortly before however requires 30 day voluntary notification of Frontier Models prior to release but does not differentiate between open and closed models.

1.1.3 The EU Cyber Resilience Act in depth

The CRA is the one of the most consequential developments and the central case study used in this Report. Its first draft pushed liability onto foundations and volunteer maintainers even if they neither sell nor profit from their code. An eighteen-month corrective effort - led by the Linux Foundation, Eclipse, Rust, Python and Apache, with GitHub marshalling pushback - produced the 'Open Source Software Steward' (Steward) as a category. Thisis the first recognition of non-commercial open source organisations in any major legislation in any jurisdiction. Concerns remain for the ecosystem where many who distribute their code at no cost might yet be caught by the regulation due to their commercialisation which is defined more broadly than royalty charges for licensing.

Drawing on thought leadership contributions from Mirko Boehm (LF Europe) and Rebecca Rumbul (Rust Foundation), this section explains why regulators misunderstand software by treating a fluid, continuously updated service as a fixed manufactured good and why the steward definition, though pivotal, remains untested until its standards and enforcement are settled.

1.1.4 AI, the Mythos Effect and Open Source under strain

Much of the report addresses how AI is altering both productivity and risk. The Mythos Effect - frontier models trained on the entirety of available code - has collapsed the time between latent vulnerability and exploit, surfacing thousands of critical CVEs and overwhelming the patching capacity of even well-resourced projects. Contributors warn of code bifurcating at scale, and of SBOMs straining as AI-generated contributions lose their provenance. The report critically examines defensive over-reactions, including sovereign forks and the NHS's decision to close-source its repositories, arguing that withdrawal from the commons negates the collaborative security benefits open source provides.

1.1.5 SBOMs, provenance and the engineering response

Software Bills of Materials (SBOMs) are both a regulatory expectation and an operational necessity. The report argues that static, checkbox SBOMs are already obsolete; effective practice requires build-time generation, enrichment with exploitability (VEX) data, and verifiable provenance through tooling such as SLSA and Sigstore. It identifies a widening UK gap, strong international engagement but no domestic statutory mandate, and proposes procurement frameworks such as G-Cloud as the most immediate, addressable lever.

The global survey and the supporting literature. A comparative survey and overview table set out how twenty-three jurisdictions treat open source, penalties for breach, and whether each follows a product, organisational or dual model. A review of the major 2025–2026 industry reports including those from Black Duck, Sonatype, the Linux Foundation and OpenSSF corroborates the central themes: a worsening, automating threat environment; transparency shifting from best practice to mandate; and a structural sustainability crisis in how open source is maintained and funded.

1.1.6 The UK's position

The report closes on the UK Cyber Security and Resilience Bill which is an organisational instrument that extends NIS-style duties but makes no mention of open source, leaving recognition to future secondary legislation. The recurring conclusion across every section is that the most valuable thing any jurisdiction can do is understand open source for what it is: critical, fragile, communal infrastructure - and regulate with the ecosystem, ensuring new burdens are shared fairly rather than crushing the volunteers who sustain it.
Open source software is foundational to modern digital infrastructure, yet its security remains a largely unregulated space. Existing cybersecurity frameworks around the world - such as the EU's NIS2 Directive or India's critical information infrastructure regime - impose duties on organisations operating critical infrastructure, but rarely address open source components those organisations might depend on.

The EU Cyber Resilience Act is the exception, and the most consequential development of the period. Its first draft would have pushed legal liability onto foundations and volunteer maintainers who neither sell nor profit from the code they produce. An eighteen-month corrective effort by the Linux Foundation, Eclipse, Rust, Python, Apache and others produced the "open source software steward" category, the first time any major jurisdiction has recognised non-commercial open source in regulation. The steward definition is pivotal, but the standards that will be drafted, and what enforcement will look like under the new categories will be an important test for the law.

Elsewhere, the pressure on open source is indirect. Organisational duties under regimes like NIS2, Singapore's Cybersecurity Act, and the UK's Cyber Security and Resilience Bill might flow upstream as supply-chain requirements (SBOMs, provenance, vulnerability disclosure) without naming open source at all. The result is a de facto regulatory perimeter built around the ecosystem rather than with it.

This report sets out information valuable to those providing responses to consultations around regulation.

1.2 Introduction to Cyber Regulation and this report

Professor Amanda Brock
CEO, OpenUK and OpenHQ

In the run up to the release of this report Sir Alex Younger the longest serving Chief of MI6 died. Alex was someone who I had discussed open source and security with on several occasions. The last thing he said to me after a discussion of why open source ought to be a sensible approach for the UK, was "may the open source be with you." This was of course in his personal capacity and long after his days at MI6. I'd like to dedicate this report to his memory, and his forward thinking.

Brexit meant that where the UK's open source community engaged in regulation and policy with the Commission for decades prior, as part of the EU, they were much less involved in the EU's Cyber regulation and creation of the Cyber Resilience Act (CRA). One of the reasons OpenUK was set up in its current form in 2019/20 was the clear need for the UK to begin to have its own industry representation and policy organisation engaging with policy makers and regulators here in the UK . The initial draft of the CRA in 2021 was so far off the mark of understanding open source, that it was greeted by a global response as an existential threat to open source.

That lack of engagement from the UK did not however mean that we were not involved in discussions in the background in the CRA discussions. That experience and engagement with our EU colleagues may be relevant her in the UK, as we move to consultation on the UK's cyber regulations and over time to subordinate regulation which we expect will impact open source.

Mirko Boehm, of Linux Foundation Europe shares his experiences in this Report'Thought Leadership calling out the challenges faced by regulators in their understanding of the modern software development landscape even before the recent impact of AI. As he explains, policymakers often try to fit old school manufacturing concepts and language into dynamic software development. This was particularly apparent in the EU process where they effectively sought to timestamp risk and liability management. Whilst this works for a manufactured item, formed as a static unit by a production line and which is easily checkable and certifiable, it is not appropriate to software.

This, however, does not sit well with the dynamic, evolving nature of software. And in particular it sits counter to the methodologies of cyber security in both software and AI which are flagged in Andrew Martin's Mythos Effect contribution. It is critical today that software is updated and security vulnerabilities with known fixes, patched. Failure to implement such is likely to be negligent. This is fully explored in the summaries of panels from our Edinburgh event last week.

As a former lawyer with 20 years experience of open source and decades of policy work, I saw a few things wrong with the original CRA in that initial format that are worth flagging for learning:

- **The regulation re-characterised software as a "good" and not a service in the EU.** Whilst software is a service under legal regimes in the UK and in many other jurisdictions, the EU has shifted software to being characterised as "a good". This has the net effect of a shift in the applicable regulation. Product Liability legislation designed for physical goods would now apply to software in the EU as a consequence of this.

  I wrongly assumed this was an error in the drafting when I first read it but this shift was rather a clear and calculated decision made with intent to capture software under product liability regulation. That is not appropriate to software which is developed in a dynamic manner today, with the upstream being incorporated into the downstream and regular updates.

The content from Mirko Boehm and Rebecca Rumbul who lead code holding foundations and are "Stewards", calls out the regulators' misunderstanding of software and its development. They emphasise that software development is ongoing and fluid. Software cannot sensibly be treated like manufactured goods fixed at a moment in time. When it comes to risk and liability, attempts to push this square manufacturing and product peg into the round hole of software may be a disaster waiting to happen.

- **A single level of liability which initially covered Foundations and the open source ecosystem as a whole.** This has evolved across the versions as a concept from the first draft but I remain personally concerned for individuals, small businesses and innovators who monetise the code they distribute free of charge and who it appears must comply under the commercialisation provisions should they have any income generated from services in relation to the software they have created. Their liability levels are higher than code held by foundations which benefits from the concept of a Steward and reduced liability.

- **The legal shift in the basis of liability from user to creator despite no monies changing hands** Historically the liability for code sat with the user who chose to use the code and not the provider on a buyer beware basis. For liability to apply there is a need for money to change hands. Whether through paying for services as an employee or through a contracted service. Open source being freely licensed and generally royalty free, there is no royalty fee and so no consideration and liability is disclaimed in open source licences. However, the provisions on commercialisation in the EU regulation shifts this and the commercialisation provisions go beyond royalties to capture any financial recompense from the code such as the provision of associated support or services.

- **Provisions allowing code to be tested and recalled** would not have work under standard open source licences which share the code in perpetuity.

As the conversation around the CRA evolved there was a huge outcry from the open source development community and pushback. I saw the negotiated response characterised as victory, the Commission creating the Open Source Steward category in the legislation benefitting from reduced liability. The potential for individuals and innovators not charging for software being deemed to be commercialising because they sell services associated with the software they create feels real.

Delving deep into the quagmire of AI, we explore Agentic in Matt Barker's content and Andrew Martin's Mythos Effect conversation, and panel from Edinburgh last week at OpenUK's SOOCon26 event in Edinburgh clarify the challenges raised by Anthropic's ClaudeMythos model and its effect. This sees a language model trained on the entirety of the volume of code available online and has the potential to enable several outcomes that could shift security in software and infrastructure beyond recognition.

Talking to long term open source contributor and founder, Jean-Baptiste Kempf it soon became clear to me that this might well play out as:

- everyone will potentially be able to fork code and to add code to this, enabling the unthinkable bifurcation of code x1000;
- At this skill there would be potentially unmanageable vulnerabilities that could be all but impossible for anyone to guarantee; and
- Software Bill of Materials which were heavily relied upon by the US public sector and have become a critical tool in software management today - explained in Sal Kimmich's contribution - may soon be irrelevant, in an AI world. Sal would of course argue that this is not the case due to long term human intervention.

We have as always sought to be as up to date as possible and included US regulation through the latest Memo from President Trump dated 5 June focused on defence and the recent Presidential Executive Order focused on AI.

All of these, accompanied by an overview of both UK and international regulation, will support our responses to the UK Cyber Bill as it progresses from its third reading in the House of Commons on 10 June.

1.3 Software Bills of Materials and the UK's Supply Chain Blind Spot

Sal Kimmich
Security Architect and Policy Manager, OpenUK

A Software Bill of Materials (SBOM) is asa machine-readable inventory of every component within a piece of software: its dependencies, versions, provenance, and known vulnerabilities. The concept has been central to US federal procurement guidance since Executive Order 14028 in 2021, and is now functionally embedded in the EU Cyber Resilience Act, which requires manufacturers to document software components as part of conformity obligations.

In the UK, no equivalent statutory requirement exists: the Product Security and Telecommunications Infrastructure (PSTI) Act 2022 establishes baseline security requirements for connectable products but does not mandate SBOM production or disclosure, and NCSC supply chain security guidance references SBOMs only as recommended practice rather than enforceable obligation.

1.3.1 That gap is widening, and two converging pressures are making it urgent.

The checkpoint model is already obsolete

Early SBOM adoption followed a familiar compliance pattern: produce a document, submit it to satisfy an audit requirement, file it. The Black Duck 2026 OSSRA report records that the mean number of vulnerabilities per codebase doubled in a single year, and that 92% of audited codebases contain components four or more years out of date.

A static SBOM in this environment provides a misleading picture of current risk. Effective SBOM practice requires generation at build time, enrichment with Vulnerability Exploitability eXchange data to distinguish reachable vulnerabilities from theoretical ones, and direct integration with remediation workflows. The OpenSSF is explicit on this point: SBOMs are operational infrastructure, not periodic documentation.

AI-generated code is breaking the provenance model

The more immediate pressure is structural. At the same time that SBOM mandates are expanding, the code those mandates are intended to govern is becoming harder to trace. Google reported in April 2026 that 75% of new code committed to its internal repositories is now AI-generated, with provenance metadata frequently stripped when that code moves into open source projects. The SPDX community is drafting model provenance extensions, but adoption is nascent. An organisation generating a build-time SBOM for a product that incorporates AI-assisted code faces a genuine gap in what that SBOM can reliably attest.

This is not an abstract future risk. The Black Duck 2026 data documents that 76% of organisations report developers using AI coding assistants against corporate policy. Shadow AI in the development pipeline produces components whose lineage is unknown to the organisation's own security team. The SBOM cannot inventory what the organisation does not know it has shipped.

1.3.2 The CRA has set the floor for UK organisations

The EU Cyber Resilience Act does not reference SBOMs by name or define them, but its obligations are functionally equivalent to such a mandate. Manufacturers placing products with digital elements on the EU market must demonstrate build-time provenance, maintain vulnerability tracking across the defined support lifecycle, and notify ENISA of actively exploited vulnerabilities within 24 hours. There is no credible path to satisfying those requirements without a maintained, automated SBOM practice.

The CRA's extraterritorial reach is frequently underweighted in UK policy discussions. UK software vendors selling into the EU market carry the full weight of CRA conformity obligations. The proposed UK Cyber Security and Resilience Bill extends supply chain obligations to operators of essential services and newly in-scope managed service providers, but introduces no equivalent product-level SBOM requirement.

UK vendors with EU exposure are already adapting to CRA-aligned standards; those operating solely in the domestic regulated market face softer obligations arriving through procurement contracts rather than statute. The practical result is competitive distortion and inconsistent visibility across shared supply chains.

Open source and the limits of provenance

SBOM obligations encounter a structural constraint when applied to supply chains built substantially on open source components. The Linux Foundation Census III found that among the most widely used open source libraries, 17% have a single developer responsible for more than 80% of commits, and a significant proportion have no formal security reporting process or signed release artefacts. A regulated organisation's ability to obtain verified provenance from a commercial supplier depends on what that supplier can obtain from its upstream. In substantial portions of the open source ecosystem, that answer is currently limited.

The CRA's open source software steward category is the first attempt by any major jurisdiction to address this in law. Stewards carry lighter, bounded obligations and are exempt from administrative fines. Liability for CRA compliance remains with commercial manufacturers placing products on the EU market. The UK Cyber Security and Resilience Bill provides no equivalent foundation. Maintainers of open source components embedded in critical national infrastructure have no direct recognition within the legislation, and the indirect pressure flowing to them from regulated entities will intensify as the Bill takes effect.

What the UK has and what it is missing

The UK policy field is not empty. The NCSC's Software Security Code of Practice provides voluntary guidance on supply chain security, and the NCSC co-authored the G7 SBOM for AI framework published in May 2026, alongside CISA, Germany's BSI, France's ANSSI, and others. That framework establishes baseline recommendations for AI supply chain transparency across seven clusters including model provenance, dataset properties, and security properties. Its co-authorship signals UK engagement with the frontier of SBOM practice at the international level.

The gap is between that international engagement and domestic statutory implementation. Voluntary guidance does not create consistent market behaviour. The evidence supports this: organisations that validate supplier SBOMs remediate critical vulnerabilities at substantially higher rates and within shorter timeframes than those that do not, but validation rates remain low in the absence of mandatory requirements.

The most immediate concrete lever is procurement. G-Cloud 15, the UK government's flagship cloud procurement framework valued at £14 billion and running to 2030, now mandates Cyber Essentials for all suppliers as a supply chain security measure. It does not require SBOMs. That gap is specific and addressable at the next framework refresh.

Engineering against current limitations

The operational reality of attested SBOM practice is harder than the policy frameworks assume, because the vulnerability infrastructure those SBOMs depend on is under acute stress. CVE submissions to the National Vulnerability Database increased 263% between 2020 and 2025, and even processing nearly 42,000 CVEs in 2025 proved insufficient to keep pace. In April 2026, NIST formally switched to enriching only the highest-priority CVEs, abandoning routine enrichment of earlier submissions. The severity scores and exploitability context that security teams depend on to act on SBOM data are now selectively available. The EU launched its own Global CVE Allocation System in response, but without enforceable interoperability commitments, defenders may need to query multiple incompatible systems to establish whether a given component is vulnerable. An SBOM enriched against fragmented, partially maintained vulnerability databases is not the risk management instrument the CRA's 24-hour notification obligations require.

The second constraint is structural. An SBOM attests to what a build contains, not to whether the process that produced it was compromised. After the GhostAction supply chain attack in 2025, provenance verification is no longer optional. The standards to address this are mature: SLSA defines build integrity requirements across progressive levels; Sigstore provides keyless cryptographic signing through Cosign, Fulcio, and Rekor; and the in-toto attestation framework binds SBOM content to verified build provenance in a structure that is tamper-evident across the full pipeline. For teams using GitHub Actions or GitLab CI, implementing SLSA Level 2 provenance generation and Sigstore signing is achievable in one to two days of engineering effort.

The implementation sequence for regulated organisations follows from this. Generate SBOMs at build time using Syft against container images and source trees, in CycloneDX or SPDX format. Sign the artefact and SBOM together using Cosign with OIDC-based keyless signing, logging to Rekor so the attestation is independently verifiable without key management overhead. Enrich against multiple vulnerability feeds rather than NVD alone, drawing on OSV, the GitHub Advisory Database, and GCVE, and attach VEX assertions to distinguish exploitable findings from theoretical ones. That final step is what transforms a static inventory into the prioritisation instrument a 24-hour remediation obligation actually demands.

The tooling is production-ready and open source. The constraint is organisational, not technical.

The priorities are clear for UK policy makers:

- First, introduce SBOM requirements into G-Cloud 15 and equivalent Crown Commercial Service frameworks at the next available refresh point, creating a de facto market standard through procurement rather than waiting for primary legislation.
- Second, develop SBOM guidance that is explicitly coherent with CRA requirements, reducing the compliance burden for UK businesses operating across both markets.
- Third, extend that guidance to address AI-generated code provenance, building on the NCSC's existing co-authorship of the G7 SBOM for AI framework rather than treating it as a separate workstream.
- And in respect of regulated organisations, the practical path is to treat CRA-aligned SBOM practice as the operative standard now. The organisations that responded most effectively to the post-Mythos CVE surge were those that already had automated dependency visibility. That capability is not contingent on domestic legislation. It is available, the tooling is mature, and the cost of not having it is no longer theoretical.

1.4 AI Agents and Identity: The Next Evolution of Enterprise Trust

Matt Barker
CEO, BoltMCP

Identity has continually evolved alongside infrastructure. In the early days of enterprise computing, organisations only needed to manage human identities. As technology evolved, we introduced physical servers, virtual machines, cloud workloads and distributed applications. Each shift required new approaches to identity, authentication and trust.

Today we face a fundamentally different challenge: AI agents that can act on behalf of humans or operate autonomously, interacting with systems originally designed for human users while consuming and acting upon data at machine speed. At the same time, the number of non-human identities continues to grow rapidly. Organisations have spent the last decade learning how to manage identities for workloads, services and infrastructure. AI agents represent the next stage of this evolution: identities capable not only of authenticating, but also of making decisions and taking actions.

We were already struggling to manage the growth of workload identities. AI agents dramatically accelerate these challenges. While AI models themselves are becoming better understood, their behaviour remains non-deterministic. As organisations expose their systems to AI, enterprise AI increasingly becomes an identity problem - do we trust the agent's identity, do we trust what it can access, do we trust what it does once it has access? This article explores these challenges.

> **Hackers don't break into your business anymore. They just find your keys in a leakage and log in.**
>
> **Matt Barker**
> **CEO, BoltMCP**

Agent Identity

AI has the potential to deliver extraordinary gains in productivity, automation and business outcomes. However, the same capabilities that make AI powerful also amplify risk.

Even before the rise of AI agents, secret sprawl had become a significant security challenge. Millions of credentials are exposed publicly every year, creating opportunities for attackers to gain access to systems and data. AI agents increase the potential blast radius by enabling credentials to be discovered, reused and acted upon at scale and speed.

Recent incidents involving exposed credentials and AI-assisted attacks demonstrate that this is no longer a theoretical concern. As agents become more capable and autonomous, establishing trust in their identity becomes increasingly important.

The industry has already learned that passwords and long-lived credentials do not scale for workload identities. The same lesson applies to AI agents. Agentic systems will increasingly require strong, cryptographically verifiable identities rather than shared secrets.

Open standards are likely to play an important role in this transition. Technologies such as SPIFFE have already demonstrated how workload identities can be established and verified at scale. These approaches may prove foundational as organisations seek to establish trust in autonomous systems.

> **Previously you could have a person in your workforce with a high level of privilege and access, and one day they might flip... Agentic makes that massively apparent to your leadership.**
>
> **Sal Kimmich**
> **Policy Manager, OpenUK**

Trust Boundaries

Establishing an agent's identity is only the first step. Once an agent is trusted, organisations must determine what it should be allowed to access.

To realise the full value of enterprise AI, organisations need to expose data, applications and business processes to AI systems. Agents are ultimately only as capable as the systems they can interact with.

There are now many ways for agents to access enterprise systems, from direct API integrations and command-line tooling through to emerging protocols such as Model Context Protocol (MCP) and other agent frameworks. Engineers are rapidly discovering new ways to connect systems and expose context.

For security leaders, however, the challenge is less about access and more about control.

How should organisations determine what an agent is permitted to access?

How should permissions be enforced as agents move across systems and workflows?

How can organisations ensure that agents remain within clearly defined trust boundaries?
It is important to adopt a consistent and centralised approach to exposing enterprise systems to AI. Without this, every integration becomes a new trust boundary with its own authentication model, permissions and operational controls. A consistent access layer provides the foundation for shared authentication, observability, governance and policy enforcement across the organisation.

Context Control

The next major challenge is not authentication, but context governance. Context is the information and tooling an AI model uses to complete a task, from documents and reports to APIs and SaaS integrations.

Most enterprise systems are designed to determine whether a user or service is permitted to make a request. Far fewer are designed to determine what context should be returned, under what circumstances, and for what purpose. This distinction becomes increasingly important in AI-driven environments.

Human users are naturally constrained by time, attention and organisational norms. Even when excessive access exists, there are practical limits to how much information a person can consume or act upon. AI agents do not share those constraints. They can traverse vast amounts of enterprise context, correlate information across systems and operate continuously at high speed.

As a result, organisations must move beyond simply authenticating requests and begin governing the context that is returned.

Can an agent access customer data but not compensation information? Can it retrieve financial information for analysis without exposing personally identifiable information? Can it use context for one purpose but not another?

These questions require a more fine-grained approach to authorisation than many organisations currently possess.

The challenge is compounded by scale. Agents can be replicated rapidly, operate continuously and interact with large volumes of enterprise data simultaneously. Fine-grained, deterministic authorisation of context will become increasingly important as organisations seek to balance AI innovation with security and governance requirements.

A pragmatic approach is to begin with human approval and oversight, granting agents limited permissions initially and gradually increasing autonomy as confidence and governance mature. As with previous generations of infrastructure, trust must be earned before privilege is expanded.

The Challenge ahead

The challenge facing organisations is not whether to connect AI to enterprise systems, but how to do so safely.

Identity, authentication, authorisation and governance are not new disciplines. What is new is the scale, speed and autonomy with which AI systems operate. Existing challenges around trust, access and credential management are amplified significantly when autonomous systems are introduced.

The organisations that succeed will be those that establish trusted identity foundations and build secure, governable context layers between agents and enterprise systems.

Those that can expose context while maintaining control over access and usage will be best positioned to realise the productivity, innovation and competitive advantages that enterprise AI promises to deliver.

You can watch these State of Open Con sessions

We have seen agents deleting databases - and after it does, you can't reverse it back.

Chaamini Mangaleswaran
Solutions Engineer, WSO2

When a human uses those permissions, they make judgment decisions... they do things that aren't likely to get them fired or destroy the company. When you give your permissions to an agent, you're allowing that LLM to have agency to use those permissions in whatever ways it sees fit - which is why over-scope permissions are so dangerous.

Matt Jarvis
Director of Applied AI, SNYK

1.5 Agentic Identity and Security Panel SOOCon 26, 5 June

Matt Barker, CEO & Co-Founder, BoltMCP
Chaamini Mangaleswaran, Solutions Engineer, WSO2
Matt Jarvis, Director of Applied AI, Snyk
Sal Kimmich, CEO, GadflyAI

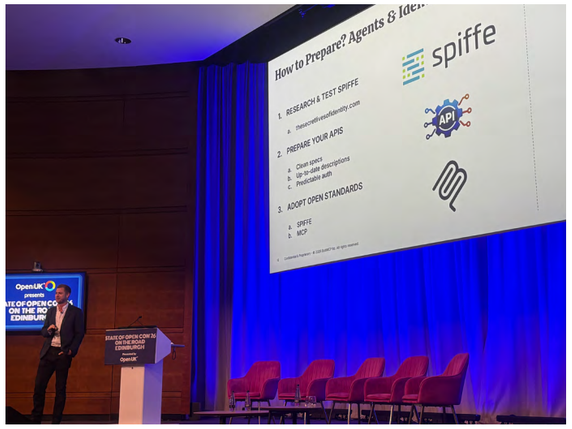

This panel examined how AI agents upend established models of workload identity and access. The recurring theme was that agentic systems make a long-standing problem suddenly acute. Organisations have always over-provisioned permissions, but humans exercise judgment about how to use them, whereas an agent handed the same permissions will use them however it sees fit. Matt Jarvis framed this as the core danger - passing your personal permission set to an LLM is highly problematic, and over-scoped permissions become genuinely hazardous once an autonomous system, not a person, is wielding them. Chaamini Mangaleswaran reinforced the point with concrete failure modes, including agents deleting databases irreversibly, and argued enterprises need a governance layer - gateways for model and MCP access - before putting agents into production.

Sal Kimmich made the case for zero trust as an enabler rather than a brake, arguing that binding permissions to a specific moment in runtime, scoped down to the kernel syscalls a workload actually needs, removes a large share of vulnerabilities and the classic 'Jekyll and Hyde' insider risk. They stressed that security is only as good as its audit trail, and that a policy document means nothing without evidence an auditor can verify.

Matt Barker pitched SPIFFE/SPIRE as the maturing standard for non-human identity - long hard to adopt, now having its 'Kubernetes moment' because agentic use cases are being built from scratch, giving teams a rare chance to design identity into the system from day one.

The panel closed on pragmatism: check your auditability before implementing AI, keep humans in the loop while trust is earned, and, in Jarvis's words, avoid magical thinking - securing AI is largely the discipline teams already know, applied at greater scale.

You can watch these State of Open Con sessions

1.6 The Mythos Effect

Andrew Martin
CEO, ControlPlane and CISO, OpenUK

Mythos didn't break open source security. It proved security was underfunded all along.

Open source is at an inflection point. We risk losing control of the greatest meritocracy in existence to actors who, out of fear or self-interest, are choosing to withdraw from the commons rather than invest in it. Open source needs support more than ever.

Anthropic's Claude Mythos Language Model

Anthropic's Mythos is a frontier AI model skilled at offensive security. It has collapsed the time between latent vulnerability and time to exploit, by ingesting entire codebases and chaining weaknesses across them at machine speed.

Mythos requires significant compute and memory: loading many lines of code into its context is what gives it the depth to unearth complex vulnerabilities across a system. In one month alone, it surfaced more than 10,000 critical software vulnerabilities. This coined the term "vulnpocalypse".

Those CVEs are not a problem for up-to-date software, good patch hygiene, and automated update pipelines. Users with rapid remediation flows are well positioned to continue to strengthen their security. But for open source, even leading projects have struggled to triage the deluge of AI-reported CVEs in the last six months, and that threatens to get worse before it gets better (encouraging signs from the Curl project suggest that it does abate).

> **The patching window has become negative. We cannot patch in time.**
>
> **Liz Rice**
> **Chief Open Source Officer, Isovalent at Cisco**

However for old software languishing in a fork, or without maintainer updates for many years, Mythos risks splintering the heart of open source itself.

The Relevance of the CRA

Open source liability, and maintainer remediation questions, have been outstanding for many years. The CRA starts to move us in a positive direction on liability and patch maintenance. But without fair maintainer remuneration, we now risk significant organisations taking matters into their own hands.

The CRA's mechanism is the reclassification of software from a service to a product, overturning the decades of "buyer beware".Instead, the first commercial distributor distributing into the EU becomes liable.

And it does not arrive alone: the EU AI Act demands transparency and human oversight of the systems we deploy. DORA and NIS2 tighten the screws on financial services and critical infrastructure. For UK businesses

serving EU customers, post-Brexit doesn't mean post-CRA: extraterritorial reach is equivalent to GDPR. The unfunded-maintainer tightrope, where hobbyists produced trillions of dollars of value with cost of failure unrepresented on balance sheets, has frayed and snapped. Mythos' effects and the CRA's remediations are both downstream of the same truth.

Sovereign Forks

Nation-states have begun positioning sovereign forks as a security posture. Their logic: open source is "too dangerous to consume unmodified", and therefore must be controlled, filtered, or replaced by a trusted national intermediary. As debated at FOSDEM 2026, the instinct is understandable and the threat is real. But the EU AI Act demands transparency and human oversight of the systems we deploy. Organisations must be able to understand what they operate.

If there is no maintainer reading the changes going into a closed fork, the requirement of understanding is a fiction: it is legislating for transparency in AI outputs, while removing the only humans capable of providing insight into the underlying code, its philosophy of design, the maintenance requirements of interface sustainability. For sufficiently complex software, the two postures cannot coexist.

The NHS withdraws

This mentality continued with the UK's National Health Service believing open source made software foundationally insecure and suggesting "temporarily" closed-sourcing everything they owned unless specifically exempted (NHS to close-source hundreds of GitHub repos over AI, security concerns).

> **Public repositories materially increase the risk... particularly given rapid advancements in AI models capable of large-scale code ingestion, inference, and reasoning (e.g. developments such as the Mythos model).**
>
> **Andrew Martin**
> **CEO, ControlPlane and CISO, OpenUK**

"NHS internal guidance"

Closed source does not protect from Mythos-level AI. Its pattern matching and blind injection will only improve, and denying the commons collaborative opportunities funnels all security inspection through a single channel, negating the well-understood benefits of open source security.

The Times are Changing

And now it continues. Organisations repackaging open source suggest they should be the arbiters of open source's safety, forking trillions of dollars of value created in the world's greatest meritocracy into a capital-generating walled garden. These "clearing house" solutions, as outlined in June's Executive Order, differentiate themselves in one fundamental way: are the patches publicly available, and upstreamed back to maintainers?

If not, then the health of open source innovation is under threat from a world of continuous, unmanaged vibe coding. Codebases will diverge on closed forks, and communities will falter.

We exist at a deeply transformative time with an incredibly rapid trajectory of change. As with all global shocks, human systems are resilient and normalisation occurs quickly afterwards.

The risk in this interstitial moment is that we accept the shock and modify the foundations of open source software. Within 6-9 months all software development will undergo the same rigorous Mythos-level testing with open source models, and open capabilities will align again.

Organisations must adapt for these tools, and do so quickly. But Mythos will not fix bad DevSecOps.

The velocity of third party software ingestion is critical. Security, speed, and the excision of vulnerable code that materialises as zero-days.

The level of test quality and configuration management comes second.

And the understanding of the systems we design, the threats and countermeasures applicable to their data classifications, and the sociotechnical interactions of humans underpin the security of even the best "secured" system.

And as organisations generate their own software at machine speed, automated remediation must also fit into those pipelines. Pipelines underpinned by the trust and security of internal systems, pipelines, and people.

The enablement of this sovereign patching and ingestion capability restores the trust in open source, equips organisations with the zero-vendor-lock-in promise of open source, and ensures the contribution of patches back upstream for the shared benefit of users in the Commons. We will reach the point of capacity for all users to run this tooling across open source, and we must be confident that this shock will become the new normal.

Liability for all

No organisation is big enough to assume the liability for all of open source, and Mythos is just the beginning of offensive models' abilities. Sovereign capacity for open source models at Mythos capability is close, and Cloudflare, OpenMythos, and automated pentest organisations have proven that the harness is the capability, not the model.

The answer is not a temporary posture. It is a permanent commitment to funding the humans who built and understand the systems we all depend on. And the contentious CRA regulation can not come fast enough.

Vulnpocalypse

The CRA isn't only regulation of software. It is the formal admission that software was never a service, and the "buyer beware" concept was always a fiction. Open source produced trillions in value while the cost of failure sat on nobody's balance sheet.

The CRA prices in that externality: it is contentious because it makes someone pay a bill that was always overdue, and the question is no longer whether, but who. Funding your upstream maintainer is no longer altruism. It is the cheapest path to defensible CRA due diligence available.

- If you're a CISO at a bank, fund one upstream maintainer whose library sits in your critical path. Not a donation, a CRA-ready contract.
- If you're a CTO, donate to the three open source projects your platform depends on most, for them to spend as they need.
- If you're a regulator, recognise that corporate sponsorship of upstream maintenance is the lowest-cost form of CRA and EU AI Act compliance available, and shape the incentives so it happens before the months-long window of extreme, AI-assisted "vulnpocalypse" pain closes in on everybody at once.

1.7 The UK Public Sector and National Health Service Response to Mythos

Clare Schramm
Technology Platforms MD, Lloyds

The UK's public services exist in a global cyber security regulatory environment that was not designed with open source in mind, yet cannot function without it. The software that runs hospitals, facilitates complex imaging, orders prescriptions, and handles clinical data is underpinned by open source components layered with proprietary systems and local configurations.

Anthropic's Claude Mythos Model

When Claude Mythos emerged, the NHS, as the largest single holder of structured, machine-readable personal data in the UK, had to quickly respond. It is unsurprising that an organisation, characterised by managing labyrinthine dependencies in the context of responsibilities that were not clearly defined by either existing or proposed regulatory frameworks in the UK, chose a maximal "risk mitigation" approach in their response.

In an environment governed by NIS Regulations 2018, the NHS issued SDLC-8 guidance on 29 April mandating that hundreds of internal open source repositories be made private by 11 May 2026.

The NHS's decision to shut down its public GitHub presence in response to the release of Claude Mythos was understandable in light of the great obligation they hold as stewards of the country's personal health information and in the context of the wider risk-based regulatory landscape in which they operate.

But ultimately, it was the wrong lever to pull

The NHS treated open source itself as a risk and the sole attack vector, rather than focussing on other attack vectors such as secrets management or governance around use of AI in the software development lifecycle. In a world where cyber regulations are evolving to include visibility and accountability across the software supply chain (as in the EU CRA), closing public repositories appears a reactive step backwards into security-by-obscurity rather than a move towards cyber resilience in a post-Mythos world.

We would much rather see a proactive response to the attack vectors alluded to in the Claude Mythos initial press release, perhaps including implementation of scanning for credentials, guardrails on use of AI tools, and temporary restriction of specific sensitive projects rather than blanket decrees to make repositories private. Such a proactive approach would align with the trajectory of regimes like the EU CRA and US SBOM practice, which operate on transparency principles that secure what an organisation can see and govern. If the UK public sector responds to each new AI threat by entirely switching off their open source presence, they risk not only eroding trust with the open source community whose code they depend on, but also risk drifting out of step with the global landscape on cyber regulations.

Cyber Regulation

Across this report, we identify that whilst global cyber security regulations have evolved, only one has established a foundational regulatory position toward open source. The CRA is the first attempt to recognise open source stewards as a distinct legal category by acknowledging that there are organisations whose role is to sustain non-commercial open source projects that are embedded in commercial products. For organisations like the NHS, who rely on many open source community contributions, they will inherently benefit from progressive regional regulations such as the CRA and should see an improvement in the security posture of upstream components over time, as stewards adopt formal policies around vulnerability handling for example. In practice, NHS partners and suppliers who also sell into the EU market will have to design their products and processes with CRA obligations in mind, and those improvements will flow into the UK market regardless of whether we adopt a similar regime.

The UK's regulatory path continues to emphasise risk mitigation for now. The proposed Cyber Security and Resilience Bill is expected to build on NIS by widening its scope to managed service providers and data centres, but we have yet to hear whether it introduces an analogue to the CRA's open source steward model. The open source libraries and frameworks that the NHS operationally relies on are high-quality, widely tested components that are maintained by communities with deep expertise, not simply an attack vector to be mitigated via contracts, scanning utilities, and procurement policies. Introducing the open source steward concept in regulation would provide standard guidance across the UK public sector resulting in fewer "knee-jerk" reactionary responses to perceived threats in the future. Public sector organisations could further consider supporting the security of the open source community they benefit from, via sponsorship of key upstream projects, contributing to security working groups, or even seconding staff into foundations where their skills can contribute to hardening and evolving key projects.

The NHS response to Claude Mythos is a case study in how a large, complex public sector organisation adapts to a world where emerging threats outpace the guardrails provided by regulations. Whilst regulation cannot secure the software supply chain on its own, it can create both the conditions and incentives for better security practices, and ensure liability sits with those who have the resources to respond.

Our hope is that the UK public sector learns from the Mythos reaction, examines their regulatory choices in light of the reliance and pervasiveness of open source in the public sector, and shifts regulatory compliance from a burden into a driver for more sustainable security practices across the supply chain.

Government Digital Services Blog Guidance

As a consequence of the NHS activity, alarm bells were ringing in other parts of the UK's civil service. The UK's somewhat dated 2011 Open Source first public sector policy is still in place and enables a considerable volume of open source work. GDS recently issued guidance "AI, open code and vulnerability risk in the public sector." The GDS guidance is explicit that AI-accelerated vulnerability discovery is a reason to raise the bar on basic operational hygiene, not to abandon "open by default" as the organising principle for public sector code.

Fortunately this repudiates the NHS stance on open source and public repositories without directly referring to the NHS. It calls out that "private repositories can create a false sense of security" and goes on to say that making code private is not an appropriate mitigation for lack of ownership, patching capability or operations assurance, and points out that systems that cannot be safely maintained should be remediated or retired. This is a firm rebuke of any "flip the GitHub visibility switch" instinct by stating that private repositories can create a false sense of security and reiterating that closing code after years in the open rarely removes real exposure, and that treating privacy as a compensating control is a red flag for under-resourced maintenance rather than a serious cyber strategy. In place of panic closures, it offers a concrete minimum standard for public code including clear ownership, secure-by-design baselines, automated hygiene, credible patch SLAs, exceptions where publication would create a specific, credible route to harm, and makes closure the time-bound exception, not the new norm.

The GDS blog clearly advocates for process over panic. Far from validating the NHS's retreat, the GDS position underlines that the answer to Mythos is not fewer eyes on public sector code, but more disciplined, well-resourced, and transparently governed open source practice across government. This validates the OpenUK proposal for a National Foundation which would incorporate many of the services that a national OSPO (Open Source Program Office) might otherwise have included. This type of collaborative, stewardship model that OpenUK has argued is essential if the UK is to remain both secure and aligned with global norms around open digital infrastructure. A need for this expertise supporting the breadth of the public sector becomes clearer by the day.

1.8 Mythos Effect Keynotes and Fireside Chat: SOOCon26, 5 June

Liz Rice, Chief Open Source Officer, Isovalent at Cisco
Andrew Martin, CEO, ControlPlane and CISO, OpenUK
Adrian Mouat, DevRel Engineer, Chainguard

The Mythos large language model from Anthropic has caused a great deal of fear in terms of software security, and this was explored in keynotes and a fireside chat during “The Mythos Effect” Session at SOOCon26, in Edinburgh on 5 June. It included keynotes and a Fireside chat.

> **Mythos is a major leap in "autonomous exploit development". It's not about whether the model can find code vulnerabilities, but whether it can exploit them.**
>
> **Adrian Mouat**
> **DevRel Engineer, Chainguard**

Three deep experts who are immersed in Project Glasswing - the industry response to this with organisations have been given access to Mythos - debated whether Anthropic’s vulnerability-finding model is brilliant marketing or an existential threat to critical infrastructure, and concluded it is both.

> **The major leap is in 'autonomous exploit development'. Presumably within Project Glasswing there are 1000s more vulnerabilities that we're never going to hear about.**
>
> **Adrian Mouat**
> **DevRel Engineer, Chainguard**

The marketing, they agreed, had at least focused boardroom attention and budgets.

> **The incredible piece of marketing has made CEOs of Fortune 100 companies incredibly aware of the risk... it's focusing the mind and budgets on: what do we do to address this?**
>
> **Liz Rice**
> **Chief Open Source Officer, Isovalent at Cisco**

The real danger is not a leap in model capability but a dramatically widened attack surface. The latest well-maintained code (curl was cited) holds up well, but five- and ten-year-old open source - often copied into firewalls and proprietary systems with no CVE correlation - is now trivially exploitable, and "the bill is coming."

> **It's not so much a huge step in model capability, but it's really just widened the attack surface to a point that... you've got ten-year-old snippets of open source code stuck in firewalls that we probably all have in our organizations. Very poor practice - and now the bill is coming.**
>
> **Andrew Martin**
> **CEO, ControlPlane and CISO, OpenUK**

Much of the discussion turned on whether supply chain security can solve this. Andrew Martin pointed to VEX (vulnerability exploitability exchange) as a stopgap for documenting mitigations and running known-vulnerable code safely, warning that the patching window has effectively gone negative because unpaid maintainers accept only a fraction of incoming patches.

Liz Rice argued supply chain hygiene alone cannot keep pace, advocating runtime mitigation via eBPF shields that neutralise an exploit in the kernel while a proper patch rolls out. Adrian Mouat made the counter-case for relentless freshness: Chainguard proved zero-CVE container images are achievable by keeping everything small and up to date, citing Google's "build horizon" of constant rebuilds. The consensus was that these approaches are complementary, not competing.

> **The major leap is in 'autonomous exploit development'. Presumably within Project Glasswing there are 1000s more vulnerabilities that we're never going to hear about.**
>
> **Adrian Mouat**
> **DevRel Engineer, Chainguard**

> **The speed of ingestion [of patches] is really the differentiator for security. Some patches do not have CVEs, they are just patched and pushed. And if CVEs are a board level metric, if that metric is not red, then what is the incentive?**
>
> **Andrew Martin**
> **CEO, ControlPlane and CISO, OpenUK**

The panel also weighed the CRA's shift of product liability onto software producers and distributors, the legal risk of "vibe-coded" patches (including a near-miss over a HashiCorp license), and maintainer burnout from AI-generated slop reports. AI, they noted, is good at finding vulnerabilities but poor at fixing them.

Closing thoughts: pay maintainers, mitigate and update - and the future is genuinely uncertain.

What 5 words sum the panel up?

"Pay open source maintainers yesterday."
Andrew Martin CEO, ControlPlane and CISO, OpenUK

"We don't know the future. It's scary."
Adrian Mouat DevRel Engineer, Chainguard

"Mitigate and update."
Liz Rice Chief Open Source Officer, Isovalent at Cisco

You can watch these State of Open Con sessions

2. Open Source, Cyber Regulation and The EU Cyber Resilience Act

2.1 The Open Source Security Landscape Globally

Open source software is a crucial part of the modern computing stack. It underpins everything from operating systems and web servers to cryptographic libraries that secure online transactions. Yet, not a few weeks go by without news of a supply-chain attack targeting open source through compromised packages, hijacked maintainer accounts, or malicious code slipped into a widely trusted repository. These incidents have repeatedly been treated as edge cases, instead of being addressed as a structural vulnerability at the heart of how modern software is built and consumed.

The security of open source software is, in a very real sense, responsible for the security of the Internet. Given this framing, one might expect regulatory interventions around the world to have adapted to be able to address open source security, however, existing regulations have yet to arrive at a satisfactory answer for how open source software should be treated.

2.1.1 Existing regulatory models

Most jurisdictions follow an organisation-centric critical infrastructure model for regulating cyber security. This includes the Network and Information Systems Regulations, 2018, the proposed UK Cyber Security and Resilience (Network and Information Systems) Bill, the EU Network and Information Security (NIS)2 Directive, Australia's Security of Critical Infrastructure Act, Singapore's Cybersecurity Act, South Korea's critical infrastructure law, Ghana's Cybersecurity Act, Kenya's CII framework, Nigeria's CNII Order, Russia's CII law, Vietnam's Cybersecurity Law, and similar regimes in the UAE, Rwanda, Thailand, and India. These laws typically impose duties around risk management, incident reporting, resilience planning, regulator engagement, and sometimes supply-chain security. However, these regimes rarely contain provisions specific to open source.

2.1.2 Regulating open source security

The most prominent attempt to bring open source software within the regulatory perimeter is the EU CRA. The CRA tries to carve out "open source software stewards" to capture organisations that support the development of open source products intended for commercial use where the organisations supporting development don't place those products on the market themselves. Stewards face more limited and targeted obligations than commercial manufacturers of software, including adopting cybersecurity policies, cooperating with market surveillance authorities, and reporting actively exploited vulnerabilities. These obligations are distinct from those imposed on manufacturers, and stewards are not subject to administrative fines for infringement of the CRA in that role.

The steward approach is novel but raises its own concerns. The concept might map reasonably well onto large foundations but would not cleanly apply to most open source activity, which is made up of small projects with informal organisation, and individual maintainers committing code. Yet, such code might often end up in commercial supply chains without the knowledge or consent of the authors. The CRA also draws its line around the concept of "placing on the market," a framework inherited from EU product-safety law, but it remains unclear whether that distinction holds up when applied to software that is freely copied, forked, and redistributed.

The CRA isn't the EU's only intervention. In June 2026 the Commission published an EU Open Source Strategy that places open source at the centre of Europe's technological sovereignty and promotes European open alternatives to non-EU proprietary solutions in critical domains. Where the CRA assigns obligations, the Strategy gestures at resources, committing to the long-term maintenance and security of critical components through stewardship, dependency analysis and an Open Source Maintenance Instrument. It also concedes the structural problem the steward category alone cannot solve, that the economic value generated by open source is frequently captured outside Europe. Whether the Maintenance Instrument is funded at a level that matters, remains – like the CRA's standards – unwritten.

The United States adopts a different approach. Executive Order 14028 on improving the nation's cybersecurity and subsequent guidance from Cybersecurity and Infrastructure Security Agency (CISA) have promoted the adoption of Software Bills of Materials (SBOMs). SBOMs are machine-readable inventories of the components within a piece of software, including open source software dependencies. The logic is that organisations cannot secure what they cannot see, and SBOMs aim to make dependency chains visible. However, SBOMs address only discovery, and not governance. Knowing that a critical system depends on a thinly maintained open source library is useful, but it does not by itself ensure that the library is secure or that anyone is responsible for making it so.

A June 2026 executive order directs the creation of an AI cybersecurity clearinghouse to coordinate and deconflict vulnerability scanning, validate findings, and prioritise patch distribution in voluntary collaboration with industry and critical-infrastructure operators, and establishes a "covered frontier model" designation based on a classified assessment of a model's advanced cyber capabilities. It expressly declines to create any mandatory licensing or preclearance requirement for releasing new models. The order says nothing about open source as a category – consistent with prior US treatment – but it signals that the centre of gravity in US cyber policy is shifting from SBOM provenance toward frontier-model security.

2.2 The EU Cyber Resilience Act

The CRA sets binding cybersecurity requirements for any product with digital elements sold in the EU. Manufacturers must design for security, manage vulnerabilities across a defined support window, ship updates, and carry a CE mark. Serious breaches cost up to €15 million or 2.5% of global turnover. It entered into force on 10 December 2024 and applies in full from 11 December 2027.

2.2.1 What it covers

The CRA applies across hardware and software products connected directly or indirectly to a network, from consumer IoT to enterprise infrastructure. Manufacturers carry the principal obligations. They must build security in by design, support products for a defined period, run a vulnerability handling process, and notify European Union Agency for Cybersecurity (ENISA) of actively exploited vulnerabilities within tight timeframes. Conformity is demonstrated through the CE mark, with higher-risk product classes subject to third-party assessment.

2.2.2 What it does for open source

Non-commercial free and open source software has a specific carve-out. A separate category, the open source software steward recognises foundations and similar organisations that support development without monetising it, with lighter obligations focused on cybersecurity policy and cooperation on vulnerability handling. The full weight of compliance stays where the law intends it, that is, on commercial manufacturers placing products on the EU market, including products built on open source components. Upstream contributors are insulated while downstream vendors are not.

This is the first time EU law has carved out a dedicated regulatory space for non-commercial open source development, and the first time any major jurisdiction has done so.

2.3 The EU Cyber Resilience Act Journey

Early outrage

The European Commission published the initial CRA proposal in September 2022, broadly highlighting manufacturers and consumers. The proposal was met with heavy criticism from open source foundations concerned about the omission of open source and the regulation's impact on developers and open source ecosystems.

Foundations, volunteer maintainers, and the loose contributor networks that produce most of the world's software were nowhere in the text, which meant they risked being swept in as "manufacturers" by default, potentially liable for software they never sold or otherwise monetised. The structural problem with the proposal was easy to articulate. Downstream vendors monetised open source software, whereas upstream maintainers did not, and treating them identically would have made hosting a widely used open source project inside the EU legally untenable for many. An attempt was made to exclude open source software divorced from commercial activity, however, the term "commercial activity" itself was also very loosely defined.

In the months following the Commission's September 2022 proposal, opposition coalesced across an unusually broad cross-section of the software ecosystem. The Open Source Initiative compiled a survey of the formal responses submitted to the Commission's consultation, noting that 131 responses were submitted to the proposed text, of which 18 (representing a significant proportion of Europe's software industry) shared OSI's concerns to some degree.

The signatories were not confined to open source advocacy groups either. Alongside the Open Source Initiative, Open Forum Europe, the Document Foundation, NLnet Labs, the Python Software Foundation, and the Electronic Frontier Foundation, the list of critics included trade associations such as DIGITALEUROPE, Bitkom, and the Information Technology Industry Council, the Linux Foundation's OpenSSF, and major corporations including Microsoft, GitHub, Huawei, and Sonatype. Despite the diversity of these voices, their objections appeared to center on a single structural flaw in the draft text. The proposal's Recital 10 attempted to exempt open source software developed outside "commercial activity," but the term was left undefined and, as the OSI put it, the exception would cause extensive problems for Open Source software in practice.

GitHub warned that the scope of commercial activity risks bringing into scope activities that are not placing a product on the market per se, while Microsoft pointed to ambiguity resulting from the intersection of OSS with "commercial activity," both in the context of infrastructure and services provided to open source projects and with regard to activities that open source projects may pursue while building OSS. Sonatype cautioned that organisations maintaining free public repositories alongside paid services might be forced to either incur substantial costs and undertake significant effort to comply with the substance of the Regulation in maintaining free FOSS repositories, or shut down the public repositories for EU users entirely.

The breadth and consistency of this feedback, spanning advocacy organisations, standards bodies, hyper-scalers, and industry associations, became one of the defining inputs to the negotiations that followed, and is widely credited with prompting the eventual introduction of the "open source software steward" category and the carve-outs for non-commercial development in the final adopted text.

Changes in the legislation during its consultation

The final regulation differs from the draft in three material ways. Non-commercial FOSS is excluded from scope. The open source software steward category was created, with lighter and clearly bounded obligations. Stewards are exempt from administrative fines. The chain of accountability stays intact, as commercial manufacturers still carry compliance duties for products built on open source components, but it no longer unnecessarily spills into producers of software who didn't directly monetise their products.

Did open source win?

Open source foundations won, however provisional that may be.

The standards that will give the obligations operational meaning are still being written. How market surveillance authorities read the steward category under pressure is a contested case, and the first major breach traced upstream decides whether the concessions hold. Further, the eighteen months of corrective lobbying came at a cost in time, money and political capital that smaller foundations could not have borne alone.

The Cyber Resilience Act reclassifies software from a service into a product. Liability now shifts to the first distributor, overturning decades of precedent.

Andrew Martin
CEO ControlPlane and CISO, OpenUK

Ultimately, the Commission produced a law that needed eighteen months of correction because it started without consulting the people who built the regulated thing. Every jurisdiction now drafting equivalent legislation has the choice of repeating that mistake or not.

Has the open source community or ecosystem won? That is a very different question.

2.4 Thought leadership: Open Source, the EU Cyber Resilience Act, and the Road Ahead

Mirko Boehm
Senior Director Community Development, LF Europe

What is the open source community?

There is no single open source ecosystem. There are many smaller groups and organisations that make individual communities and which collectively form the "open source ecosystem" of individuals and organisations often referred to as "The Community". The Community has organisations which coordinate the production of code and associated standards, etc. As we explore open source it is helpful to have this understanding to give clarity.

The Linux Foundation (LF) is a code-holding foundation founded in 2000 and which has grown to be the world's largest project holder with a turnover in the hundreds of millions of dollars. It supports around 1,300 projects of varying sizes, providing both administrative and fiscal support to the communities behind them through a model of subsidiary foundations such as the Agentic AI Foundation and Cloud Native Computing.

LF spun out a European chapter, "Linux Foundation Europe" (LF Europe), in September 2022. The Foundation itself is largely neutral and apolitical but participates in policy work where LF Europe has been very active. LF Europe is funded by members and the Linux Foundation. Its members include Ericsson, Hitachi, Qualcomm, Red Hat, and others, listed here.

The Open Source Security Foundation (OpenSSF) is one of the LF's subsidiary foundations through which we work when open source software intersects with security, which publishes reports created by LF's research arm and provides guidance and open source projects for sustaining secure development and consumption of open source.

What regulators tend to miss

When regulators look at open source software, they most often misunderstand how software is developed generally and the more specific nuances of open source. They reach for the familiar categories - "Manufacturer" and "Consumer" for example. These are used for very different purposes in more static creation of goods which are fixed for delivery and may be certified. Often Regulators try to fit this language to both software and in particular open source software. The dynamic nature of software means that such a structure suitable for the production of goods, does not sit well with it. Whilst this familiar language and the concepts behind manufacture achieve the regulatory goals in traditional goods which are static and can be signed off at a moment in time at which risk is assessed and liability understood, this does not work well for software development.

The European Commission also tends to undervalue projects run under foundations like the Linux Foundation. A foundation is an established, independent organisation with formal governance and a sustainable funding model. We already have these governance and funding structures in place. It is not "just one person" maintaining software projects. Regulation success requires understanding of the differences between foundations, projects and individuals in their creation and maintenance of software.

Open Source Stewards as a Concept

Under the Cyber Resilience Act (CRA), the "open source software steward" (Steward) is introduced to regulation as a concept. It was created in response to feedback on the initial CRA proposal which referred only to "manufacturers" and consumers of open source software and failed to understand the role that foundations play in open source - part of the regulatory lack of understanding of how open source works. For the purpose of the CRA manufacturers are a natural or legal person who develops or manufactures products with digital elements or has products with digital elements designed, developed or manufactured, and markets them under its name or trademark.

Stewards are "Legal Persons" being individuals and organisations other than Manufacturers being commercial organisations or individuals, whose purpose is to provide sustained, systematic support for the development of open source software intended for commercial use, and to ensure those products remain viable. In reality, Stewards are mostly foundations like LF and LFE, that is, organisations that manage projects but do not themselves monetise the software commercially in paid distribution and support. Also we must understand that all open source software is potentially commercially used.

Stewards have the advantage under the CRA of a reduced level of liability to recognise their role in the open source ecosystem. This will not apply to open source software developed and maintained by individuals or organisations that are not considered Stewards under the CRA.

Projects without Stewards

A significant amount of critical open source software is not maintained by Foundations. It is maintained by individuals or small teams. How those projects get secured, and how regulation should treat them, is an important topic, and probably deserves its own conversation, but for now, that software will be treated as it has the potential for commercial use, as open source software that does not have the benefit of a Steward.

On the CRA journey

As the CRA stands today, I do not have many worries about it.

The initial draft was a different story however. There were real problems, and the response from parts of the open source community was alarmist. To some extent that alarmism was necessary to get the point across to policymakers during the consultation period, because the draft had been created without anyone speaking to the people building the software being regulated, or having a full understanding of what was being regulated.

Since then, the law has in my opinion improved in several concrete ways. It now largely exempts non-commercial and non-profit open source hosted by Stewards from its regulatory obligations, separates development from distribution, and draws much clearer lines around what counts as commercial activity. It also created the light-touch steward role, making the CRA the first law to explicitly distinguish commercial manufacturers from the non-profits that take responsibility for concrete open source releases. The alarmism, in some quarters, has not improved, and I do not agree with continuing to sound the alarm about a law that has materially changed.

What matters most now is what comes next: the standards being developed and implemented, and market surveillance by the regulator.

The Standards

The CRA requires that Standards are implemented to achieve a presumption of conformity with its essential cybersecurity requirements. Harmonised standards translate those essential requirements into detailed technical specifications, and a product that conforms to them benefits from a presumption of conformity with the CRA's essential requirements.

A maintained open source software library is reused across thousands of unrelated non commercial projects and commercial products. The library does not belong to any of the manufacturers of these projects or Products. A steward cannot certify such a component delivered by it against a vertical, product-specific standard for either non commercial or commercial distribution, because it does not ship the library as a product in any one category. Instead, it maintains "upstream code" that others integrate "downstream" into their projects and products which projects and products may be of many different natures and across many product categories or verticals.

A Steward can however reasonably demonstrate sound processes, secure development, coordinated disclosure, structured vulnerability handling in its creation and maintenance of the Library. These are precisely the things horizontal, product-agnostic standards can address. Standards under the CRA will need to be both vertical and horizontal, because open source itself is horizontal.

Lessons for the UK from the CRA

The mistakes of the European Commission in the initial CRA draft should not be repeated by the UK Regulators. Drafting laws and policies that touch software and open source must be a consultative process from the start and this should begin by engaging the open source communities and their representatives - The Community - in the UK, before the regulation is drafted.

The most important regulation open source has engaged with

Without question, the CRA has been the most important piece of regulation for the LF and Foundations generally in recent years. It is the file we have spent the most time on.

There is currently no other regulation we are aware of with comparable impact on open source developers, though similar activity may be emerging in global jurisdictions including Korea and Japan.

Ideally, cybersecurity regulation should provide guidance to the Community, guidance that can stay fluid and evolve over time, rather than bind it to fixed provisions it cannot realistically meet. Technology and the open source ecosystem move faster than primary legislation can be amended, so rules that are rigid on the day they pass are often outdated by the day they apply. Guidance can be revised; a statute cannot be, easily.

2.5 Fireside chat: The CRA and Standards

Rebecca Rumbul
CEO and ED, Rust Foundation

Rebecca helped build the Rust Foundation from the ground up as its CEO and Executive Director. The Rust Foundation is a non-profit that was established in 2021, and is the steward for the Rust programming language. Support for Rust's security advantages around Memory Safety went as far up as the White House in the last Administration.

1. When the Cyber Resilience Act was first being drafted, what was the open source community's reaction, and how did the Rust Foundation get involved?
When the CRA was being developed, it became very clear to a lot of us in the open source space that the Commission was only working with two definitions: "manufacturers" and "consumers". That didn't provide for the open source community or its foundations at all.

Organisations like the Rust Foundation, Eclipse Foundation, and the Python Software Foundation are neither consumers nor manufacturers. We are managing huge, sometimes very loose coalitions of individuals producing software that is eventually used by others in products that may be commercially distributed. But the coalitions of people and organisations, sometimes referred to as the "Open Source Community" are themselves not necessarily monetising the code directly or indirectly and may not be doing business with it in any way.

So a coalition of foundations - the Linux Foundation, Eclipse, the Rust Foundation, and a few others started to engage directly with the European Commission. A number of the policy organisations in the EU that understand open source got involved too.

There was a lot of activity to try and educate the Commission about software and open source and how it works. This was aimed at getting them to understand that the legislation as it was being drafted could be potentially catastrophic for open source foundations if we were determined to be "manufacturers" under the regulations, because of the kinds of obligations that would put on us and the potential liability we could incur. This in turn would have a knock on effect on the ability of the foundations to create, maintain and distribute open source software with these coalitions of individual and organisational partners.

2. The "Open Source Steward" concept in the CRA was the eventual result of that engagement. How significant is that, and where do things stand now?
The "Open Source Steward" concept is actually a huge achievement by those engaging with the Commission. This is the first time the idea of an Open Source Steward has been recognised in EU regulation or any regulation. But the Bill has now passed into law, and we're in a phase where everyone is trying to figure out what these definitions actually mean in practice, and what this means in relation to the letter of the regulation. There's a lot of talk at the moment, but very little certainty about what it means or how we can manage its requirements.

It's quite difficult to provide any certainty to our open source community right now because, ultimately, we do not know exactly how some of this legislation is going to be interpreted legally.

3. How are the foundations working together to prepare for implementation?
The Eclipse Foundation and the Linux Foundation have both set up working groups to try and figure out what the obligations are going to be for Stewards under the CRA and what we need to do as organisations classed as Stewards. It's mostly at the talking stage right now, but those working groups are producing a lot of useful material giving guidance on what the obligations are and how you might understand or interpret X, Y or Z provisions.

That should get us to a place where, when the law is finally implemented, we know roughly what we're doing.

But we are still waiting on case law to know exactly how the Commission plans to interpret the wording in the Regulation. For now the best we can do whilst we wait years for case law is interpret it.

Key Provisions EU Cyber Resilience Act

Subject	Provision
Liability under the CRA	turns on commercial activity and placing a product on the market - not on whether code sits in a foundation. Three positions follow.
Manufacturer (Art 3(13))	anyone who develops a product and markets it under their own name - individual or organisation, for-profit or not - who is first to place a monetised product on the market. Full obligations, including conformity assessment and fines.
Steward (Art 3(14))	a legal person giving sustained support to commercially-used open source without placing it on the market. Light-touch obligations (Art 24); no fines (Art 64(10)).
Out of scope	un-monetised FOSS and upstream contributors - no CRA liability even with millions of commercial users; liability rests with the manufacturer who integrates the code. The trigger is monetisation - charging, paid support beyond cost recovery, pay-with-data, or donations exceeding costs - not how the software was developed or funded.

Figure 1

4. Stepping back a bit – how has the broader conversation around cyber regulation changed within open source communities over the last few years?

Cyber Regulation and open source has definitely changed over the last five or six years.

People started to wake up to the scale of potential security issues which is unsurprising, with the adoption of code that's open source at scale in commercial software stacks. In the US, under the Biden administration, they invested heavily in the Cybersecurity and Infrastructure Security Agency (CISA). In February 2024 the White House put out a white paper that we consulted on with them while they were writing it. It promoted the use of memory-safe languages like Rust as default good practice. These are things - including SBOMs - that can reduce a lot of the bugs that find their way into software.

That work was curtailed by the current administration, so it has fallen by the wayside a little bit. That said, a lot of commercial companies took the messaging to heart. The big hyperscalers have invested heavily in trying to secure open source which forms a huge part of the commercial stack in their platforms. They have put funding into the ecosystems to enable foundations like the Rust Foundation to hire security people and start professionalising how we do security.

That has been very successful. But the kind of funding we need to do this properly and sustainably is still a lot higher than what is available globally, despite that being tens of millions of dollars.

5. What do regulators most commonly misunderstand about how open source software is built?

There's an awful lot that regulators don't really understand about how software is built these days. A lot of people just seem to assume that Microsoft writes every single bit of code for every one of its products, and that that's how it works everywhere else. Most people have no idea that this open source ecosystem exists – that it's "upstream" and that everything else "downstream" relies on that and draws from it, even the biggest commercial entities.

It's no surprise, really, that it's difficult for regulators to regulate this sector when they simply do not have the people available to them to explain how it works and to map it out.

What we've seen with bodies like the EU Commission and other governments is that they try to regulate software starting from a position of just assuming there are consumers and manufacturers and nothing in the middle – and then they have to backwards-engineer and almost fudge things a little bit to accommodate the fact that open source is a thing. It would be much more effective for these organisations to recognise there is a skill and knowledge deficit, and to get the right people in earlier on to build legislation or regulation that is actually fit for purpose.

6. Lobbying often carries a negative connotation, but open source foundations are increasingly engaging with policymakers. How has the community's approach to that work evolved?

The open source sector has been a little bit resistant to full professionalisation in the past, but I think there's a recognition that the community as it once was doesn't necessarily have the policymaking and lobbying skills we now need. We can't fly under the radar anymore. We need people with a massive range of skills who understand how these worlds work as well as how the software worlds work.

Coalitions of organisations have done a really good job over the last few years of recognising this and moving towards a more professional position in terms of engaging with policymakers and lawmakers. Eclipse and Linux are the two biggest foundations – they're the organisations that can afford to hire a policy officer to work on these things full-time. They are incredibly generous in creating working groups and bringing the rest of the foundations like us, that are a lot smaller, into that – consulting with us and pushing the work forward. This was a huge part of our being able to have the CRA recognise the Open Source Steward concept.

7. Beyond the language being memory-safe, how does the Rust Foundation contribute to security across the ecosystem?

We work in coalition with the Linux Foundation on our security hygiene and security posture. I'm on the board of the OpenSSF - the Open Source Security Foundation – to make sure we are working with other people to direct funding and prioritise security across all of the ecosystems.
We are very invested in security in Rust.

As a memory-safe language, Rust tends to be a go-to for writing new software that will avoid memory bugs, but security across ecosystems is so much more than that. We're very conscious that we need to be actively investing to ensure people can have confidence in using the language.

8. Finally - what does good regulation in this space look like to you?

I'm not against regulation.

Years ago, I used to work for a couple of different regulators. Good regulation enables innovation. Good regulation enables growth and supports it. Good regulation needs to understand what a bad actor is versus what an uneducated actor is.

And it requires that the regulators themselves have a detailed understanding of the industry they're regulating. That's sometimes what's missing – there are very good intentions, but very little understanding of how things will be distorted or affected, what additional costs or barriers it puts up.

I used to work for the Information Commissioner's Office, and we did some work on the development of GDPR. At no point during the development of that piece of regulation did anyone think that the result would be everyone in the world being annoyed by abrasive cookie banners. Those are the kinds of ripples that happen when you aren't thinking about the peripheral things.

3. Global Cyber Security Regulation and Open Source

3.1 Survey of Global Regulation

We surveyed cyber security regulation across 23 jurisdictions, finding universal legislative activity, but almost no coordinated effort on open source. Twenty-one of the 23 jurisdictions reviewed have a cyber security law enacted, in force, or in draft stages. The instruments range from sweeping product regulations like the EU's Cyber Resilience Act and the UK's Product Security and Telecommunications Infrastructure Act, to organisational duties under regimes such as NIS2, Singapore's Cybersecurity Act 2018, India's CERT-In Directions, and China's Cybersecurity Law and associated framework. Only the African Union and South Africa have no dedicated cyber security instrument in our survey.

3.1.1 Regulation Addressing open source

Of the 23 jurisdictions surveyed, only two - the EU and the US - address open source software substantively within cybersecurity regulation. A small number of others touch it more lightly and outside binding cyber law (Canada via its AI Strategy, Qatar via Q-CERT/NCSA guidance, Rwanda by implication only). Open source underpins the digital infrastructure these laws seek to protect, but in most regimes it is regulated only indirectly, through duties placed on the organisations that consume or distribute it. A hospital running Linux in the UK is regulated under NIS Regulations 2018 as an operator of essential services, but maintainers of the Linux kernel are not. A bank in Singapore using an open source cryptographic library is bound by the Cybersecurity Act 2018, but the library's contributors are not. In jurisdictions like China, Vietnam, Russia and Thailand, the regulatory weight falls on network operators, CII owners and service providers, with no carve-out, recognition or accommodation for the upstream open source ecosystem on which those operators depend.

3.1.2 The European Union (EU)

The EU CRA stands out in the survey as it is the only law that creates a dedicated regulatory category for an "open source software steward." The steward category acknowledges non-commercial development of software with lighter duties focused on cybersecurity policy and cooperation on vulnerability handling. Stewards are exempt from administrative fines, and non-commercial FOSS is generally exempt when developed and supplied outside a commercial activity. This term is broader than commercial licensing and royalty bearing and potentially captures the provision of services such as support which individuals, innovators and small organisations might supply for code they distribute free of charge.

They are classed as "Commercial manufacturers" that place products with digital elements on the EU market, including products built on open source components, carry the full weight of CRA compliance, with penalties of up to €15 million or 2.5% of worldwide turnover. Individuals and parties of limited economic standing would not enjoy the reduced liabilities available to a Steward. The result could be a bifurcation of open source models - between code hosted under a Steward or Foundation and code falling outside such arrangements. This may see a rise in the number of foundations and force codebases into foundations.

3.1.3 The United States (US)

In the run up to publication of this Report, Donald Trump issued an Executive Order requiring a 30 day review by the US of certain models and on 5 June a National Security Presidential Memorandum NSPM-11 rescinding the Biden-era NSM-25 Presidential Memo on security. The Memo specifically states that the national security enterprise should "adapt the best commercial and open source technologies for mission use in respect of open source". It does not mention space but formalises governance around the Pentagon's classified-network deals with eight firms - SpaceX which will have its initial public offering on 10 June - and sets new rules on autonomy, vendor dependency and accountability.

The US approach has been different but with a similar effect to the EU. EO 14028 and its implementing instruments (NIST's Secure Software Development Framework, SBOM guidance, federal procurement clauses) treat open source as an integral part of software supply-chain assurance for federal procurement. There is no stewardship category and no standalone OSS regulator, but the practical consequence – that suppliers must demonstrate provenance, integrity and vulnerability management across open source components – is comparable. Enforcement is through procurement contracts rather than through fines.

The US has legislation on open source and supply chain (through Executive Order 14028 and its implementing supply-chain work) and this includes SBOMs which are looked at in more detail in 1.3 by Sal Kimmich.

Other Significant Jurisdictions

Qatar does address open source through Q-CERT/NCSA guidance for using OSS, but not substantively. Rwanda and Qatar further have provisions that may apply to open source by implication, but offer no clear treatment.

3.1.4 The Focus: organisation versus product focus

Regulatory instruments from the survey can be grouped into three categories: product, organisational, and dual focused regulation.

Regulatory Focus

Focus	Overview
Product-focused regimes	are the minority but the most directly relevant to open source as code. The CRA, the UK Product Security and Telecommunications Infrastructure regime, the US Cyber Trust Mark, and the UAE's National Policy for Internet of Things Security all regulate what is sold rather than who uses it. These are the regimes most likely to affect open source ecosystems, because they reach into the software supply chain, although their obligations fall primarily on manufacturers and other commercial actors rather than directly on open source maintainers. Out of these regulations, only the CRA distinguishes commercial supply from non-commercial code contribution.
Organisation-focused regimes	are the dominant model worldwide. NIS2 in the EU, the proposed Canadian Critical Cyber Systems Protection Act, the proposed UK Cyber Security and Resilience Bill, Singapore's Cybersecurity Act, Korea's Information and Communications Infrastructure Protection Act, India's CERT-In Directions, Ghana's Cybersecurity Act 2020, Nigeria's CNII Order 2024, Vietnam's Law on Cybersecurity, and many others all impose duties on the operators of regulated infrastructure rather than on developers of underlying technologies. These regimes affect open source indirectly by dictating what regulated entities require from their suppliers and what upstream assurances they extract.
Dual regimes	cover both products and organisations and are emerging in Australia (Cyber Security Act 2024), China (Cybersecurity Law and associated framework), the UK (the PSTI Act read with the NIS Regulations) and the United States (EO 14028). These are the most likely to set the regulatory direction for the next decade, because they recognise that securing digital infrastructure requires duties on both ends of the supply chain.

Figure 2

3.1.5 Penalties for Breach

Penalties vary widely across regimes. Information about standout regimes is below.

Penalties for Breach of Cyber Regulation

Nation and Regulation	Penalty for breach
EU CRA	up to €15 million or 2.5% of worldwide annual turnover for serious breaches.
EU NIS2	up to €10 million or 2% of turnover for essential entities; up to €7 million or 1.4% for important entities.
UK PSTI	monetary penalties up to £10 million or 4% of qualifying worldwide revenue.
UK NIS Regulations 2018	up to £17 million depending on contravention.
US CIRCIA	subpoenas and civil enforcement mechanisms (rule pending finalisation).
China	fines, suspension, licence revocation, confiscation and personal liability.
Russia, Vietnam, Thailand	administrative sanctions plus criminal liability and, in Russia's case, import-substitution and procurement-exclusion consequences.

Figure 3

For open source specifically, the CRA's Steward distinction is worth highlighting.

Stewards are not subject to administrative fines at all, although entities acting as manufacturers or distributors using open source components face the full penalty regime, while shielding upstream contributors. Although the remainder of the requirements are on manufacturers and not community contributed code, the definition of commercialisation may be broad enough to catch anyone paid around their codebases. This structure may not achieve its goals and we must await clarification.

3.1.6 What this means for open source

First, the world has decided that cyber security is a matter for regulation, but it has not decided what to do about open source. The CRA's stewardship model is a serious attempt at an answer; the US procurement-led model is another; almost every other jurisdiction has deferred the question. As regulators in the UK, Canada, Australia, France, Germany and elsewhere update their frameworks, the choice they face is whether to follow the CRA's lead, adapt the US model, or continue to treat open source as a problem for someone else.

Second, the indirect pressure on open source is intensifying even where direct regulation is absent. Organisational duties under NIS2, the proposed UK Cyber Security and Resilience Bill, Singapore's Cybersecurity Act, and equivalent regimes elsewhere translate, in practice, into supply-chain requirements that flow upstream to maintainers. SBOMs, vulnerability disclosure expectations, and provenance guarantees are becoming de facto obligations even where they are not required by law.

Third, the absence of open source-specific provisions in most jurisdictions is not neutrality but neglect. The regulatory environment is being built around open source rather than with it. The CRA, for all the controversy around its initial draft, is the only major law that attempts to question what regulating open source could look like. Whether that model is replicated elsewhere will determine whether the next decade of cyber regulation works with the open source ecosystem or not.

3.2 Global Cyber Regulation Overview Table

Jurisdiction	Cyber law?	Name, year and Link	Status	Any ancillary regulation or associated standards? If yes, what?	What is regulated?	Does regulation cover open source?	Any open source-specific provisions and if yes, what?	What areas of open source are covered - software, hardware, data, standards or AI
Canada	Yes	Critical Cyber Systems Protection Act, proposed through Bill C-8, 2025	Draft	Would require regulations/orders designating operators, cyber security programs, supply-chain risk management and incident reporting	Organisations	No	No OSS-specific provision	N/A
Canada	No	Canada AI Strategy	In force	N/A	Both	Yes	Aims to have canada invest in and sustain open source AI development and responsible adoption	AI
UK	Yes	Cyber Security and Resilience (Network and Information Systems) Bill, 2025	Consultation	Amends Network and Information Systems Regulations 2018; would require and sector regulator guidance	Organisations	No	No OSS-specific provision	N/A
Australia	Yes	Cyber Security Act 2024	Enacted	Cyber Security (Security Standards for Smart Devices) Rules 2025; ransomware reporting rules; Cyber Incident Review Board framework	Both	No	No OSS-specific provision	N/A
Australia	Yes	Security of Critical Infrastructure Act 2018	Enacted	SOCI Rules; Risk Management Program Rules; incident reporting rules	Organisations	No	No OSS-specific provision	N/A
China	Yes	https://www.chinalaw-translate.com/en/22075-2/	In force	Multi-Level Protection Scheme (MLPS 2.0); Critical Information Infrastructure Security Protection Regulations; Cybersecurity Review Measures; Data Security Law; Personal Information Protection Law	Both	No	No OSS-specific provision	N/A

This table is for the purpose of information only and does not provide legal advice or substitute the need to get such.

Jurisdiction	Cyber law?	Name, year and Link	Status	Any ancillary regulation or associated standards? If yes, what?	What is regulated?	Does regulation cover open source?	Any open source-specific provisions and if yes, what?	What areas of open source are covered - software, hardware, data, standards or AI
EU	Yes	Regulation (EU) 2024/2847, Cyber Resilience Act, 2024	Enacted	Harmonised standards to be developed; CE marking; conformity assessment; market surveillance; ENISA vulnerability reporting	Product	Yes	Yes. Non-commercial FOSS is generally excluded where developed and supplied outside a commercial activity. However, obligations can still arise where OSS is integrated into products with digital elements or where an entity acts an "OSS Steward" (category with lighter limited duties, including maintaining a cybersecurity policy, cooperating on vulnerability-handling and disclosure, and supporting secure development practices.)	
EU	Yes	Directive (EU) 2022/2555, Network and Information Security Directive 2, 2022	Enacted	National implementing laws; ENISA guidance; sectoral rules; incident reporting standards	Organisations	No	No OSS-specific provision	N/A
India	Yes	CERT-In Directions under IT Act s.70B, 2022	In force	CERT-In FAQs; IT Act 2000; sectoral cybersecurity rules; DPDP Act for personal data	Organisations	No	No OSS-specific provision	N/A
Qatar	Yes	National Information Assurance Policy v2.0; Q-CERT OSS Guidelines	In force	NIA Policy v2.0; NIA standards; Q-CERT/NCSA guidance; sector rules e.g. financial sector cyber rules	Organisations	No	Yes, but guidance only: Q-CERT/NCSA Cyber Security Guidelines for Using Open Source Software; not a dedicated OSS statute	
Singapore	Yes	Cybersecurity Act 2018	In force	Cybersecurity (Critical Information Infrastructure) Regulations 2018; licensing rules for cybersecurity service providers	Organisations	No	No OSS-specific provision	N/A

Jurisdiction	Cyber law?	Name, year and Link	Status	Any ancillary regulation or associated standards? If yes, what?	What is regulated?	Does regulation cover open source?	Any open source-specific provisions and if yes, what?	What areas of open source are covered - software, hardware, data, standards or AI
South Korea	Yes	Act on the Protection of Information and Communications Infrastructure	In force	Information and communications network laws; PIPA; sector cybersecurity rules; KISA guidance	Organisations	No	No OSS-specific provision	N/A
UAE	Yes	Critical Information Infrastructure Protection Policy	In force	NIAF; Dubai DESC ISR/CII rules; sector-specific regulators; federal cybercrime and data laws	Organisations	No	No OSS-specific provision	N/A
UAE	Yes	National Policy for Internet of Things Security, 2023	In force	IoT security policy directives; may interact with product, telecom, data and sector rules	Product	No	No OSS-specific provision	N/A
UK	Yes	Network and Information Systems Regulations 2018 - https://www.legislation.gov.uk/uksi/2018/506/contents	In force	NCSC Cyber Assessment Framework; sector-specific competent authority guidance	Organisations	No	No OSS-specific provision	N/A
UK	Yes	Product Security and Telecommunications Infrastructure Act 2022; UK product security regime	In force	Product Security and Telecommunications Infrastructure (Security Requirements for Relevant Connectable Products) Regulations 2023; UK product security guidance	Both	No	No OSS-specific provision	N/A

Jurisdiction	Cyber law?	Name, year and Link	Status	Any ancillary regulation or associated standards? If yes, what?	What is regulated?	Does regulation cover open source?	Any open source-specific provisions and if yes, what?	What areas of open source are covered - software, hardware, data, standards or AI
USA	Yes	U.S. Cyber Trust Mark, FCC IoT cybersecurity labeling program	In force	NISTIR 8425-derived IoT cybersecurity criteria; FCC cybersecurity labeling administrator framework	Product	No	No OSS-specific provision	N/A
USA	Yes	Cyber Incident Reporting for Critical Infrastructure Act, 2022 - https://www.cisa.gov/circia	Enacted	CISA rulemaking; critical infrastructure sector definitions; incident/ransom payment reporting rules forthcoming	Organisations	No	No OSS-specific provision	N/A
USA	Yes	Executive Order 14028, Improving the Nation's Cybersecurity, 2021	In force	NIST SSDF SP 800-218; NIST SBOM guidance; OMB memoranda; federal procurement clauses	Both	Yes	Yes, but not as OSS-maintainer regulation. EO 14028 and implementing software-supply-chain work expressly include software supply-chain security and open source software as part of federal software assurance	
USA	No	Presidential Memorandum on Artificial Intelligence (AI) in the National Security Enterprise NSPM-11	In force		Both	Yes	Explicitly states open source AI technologies may be adapted for use by the national security enterprise	AI
France	Yes	NIS2 implementation / French critical-entity cybersecurity framework	Draft	ANSSI guidance; existing OIV/OSE rules; EU CRA applies directly for products	Organisations	No	None beyond EU CRA	N/A

Jurisdiction	Cyber law?	Name, year and Link	Status	Any ancillary regulation or associated standards? If yes, what?	What is regulated?	Does regulation cover open source?	Any open source-specific provisions and if yes, what?	What areas of open source are covered - software, hardware, data, standards or AI
Germany	Yes	NIS2 / BSI implementation background	In force	BSI Act, KRITIS rules, sector-specific laws; EU CRA applies directly for products	Organisations	No	None beyond EU CRA	N/A
Ghana	Yes	Cybersecurity Act, 2020, Act 1038	In force	CII directives; licensing/accreditation rules; CSA guidance	Organisations	No	No OSS-specific provision	N/A
Japan	Yes	Basic Act on Cybersecurity, 2014; METI/NCO Guidelines on Cyber Infrastructure Providers, 2026	In force	N/A	Both	No	No OSS-specific provision	N/A
Kenya	Yes	Computer Misuse and Cybercrimes Act, 2018 plus CII Regulations 2024	In force	N/A	Organisations	No	No OSS-specific provision	N/A
Nigeria	Yes	Designation and Protection of Critical National Information Infrastructure Order, 2024	In force	National Cybersecurity Fund provisions; CNII Order 2024; sectoral CBN/NITDA rules	Organisations	No	No OSS-specific provision	N/A

Jurisdiction	Cyber law?	Name, year and Link	Status	Any ancillary regulation or associated standards? If yes, what?	What is regulated?	Does regulation cover open source?	Any open source-specific provisions and if yes, what?	What areas of open source are covered - software, hardware, data, standards or AI
Russia	Yes	Federal Law No. 187-FZ on Security of Critical Information Infrastructure, 2017	In force	Import substitution rules; domestic software registry; data localisation law; sector security rules	Organisations	No	No OSS-specific provision	N/A
Rwanda	Yes	Minimum Cybersecurity Standards for Essential Service Providers 2023	In force	Application Software Security Directives; Minimum Cybersecurity Standards; data protection law 2021	Organisations	Unclear	Application software security standards may apply regardless of source model	Unclear - application security if applied to open source
African Union	No	N/A	N/A	N/A	N/A	N/A	N/A	N/A
South Africa	No	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Thailand	Yes	Thailand Cybersecurity Act B.E. 2562 (2019)	In force	CII rules, sector notifications, incident reporting and cybersecurity codes of practice	Organisations	No	No OSS-specific provision	N/A
Vietnam	Yes	Law on Cybersecurity (2018)	In force	Decree 53/2022; data localisation and local presence rules; network information security rules	Organisations	No	No OSS-specific provision	N/A

4. What the Literature Says

4.1 Overview

The current literature landscape on cybersecurity 2025 to 2026 converges on a picture of open source as critical-but-fragile infrastructure facing simultaneous pressure from three directions: a more hostile and automated threat environment, a new regulatory regime demanding accountability and transparency, and AI tools that accelerate both progress and peril. The proposed path forward is consistent across sources - automation, transparency, shared responsibility, and a rebalancing of effort from consumption toward sustainable stewardship, while carefully protecting the volunteer-driven model that makes the ecosystem viable.

4.2 The Reports

The reports collectively document a worsening vulnerability landscape amidst changing AI imperatives. The Black Duck OSSRA Software Governance in the AI Era Report from March 2026 reports vulnerabilities more than doubling (up 107%), with 87% of codebases containing at least one vulnerability and supply chain attacks hitting 65% of organisations. It also notes that the number of open source components rose by 30%, while license conflicts appeared in 68% of the codebases examined.

Sonatype's 2026 State of Software Supply Chain report frames open source malware as a 'nation-state business model' targeting developer credentials and build environments. A shared frustration is the persistence of avoidable risk. Sonatype notes known-vulnerable versions remain in use even when fixes exist, driven by workflow inertia and unclear ownership.

4.2.1 Regulation as a Driving Force

Published contemporaneously with this report, the Open Source Software Security Foundation (OpenSSF) CRA Awareness and Readiness report sets out the current challenges in realising the directives of the CRA, and notes that many do not even realise the approach of deadlines, particularly that of December 2027 as the full compliance target year. The report identifies that 66% of respondents remain unfamiliar with the CRA, similar to the level in 2025 (62%), despite the regulation entering into force during the year. 41% of organisations already familiar with the CRA have still not determined whether the regulation actually applies to them. Only 41% of manufacturers expect to be fully compliant by December 2027, while 39% do not know when they will be.

With a focus on new legal roles and liability concerns, both the paper published by Liane Colonna and the report The Linux Foundation's Pathways to Cybersecurity Best Practices in Open Source define an emerging 'open source software steward' role and clarify the distinct responsibilities of contributors, stewards, and manufacturers.

Colonna's paper gives an overview of the adoption of the CRA and its interface with the Product Liability Directive (PLD). For Colonna, this marks a major transformation in ensuring software security and accountability. The paper investigates the scope of exemptions for non-commercial open source projects while defining the separate roles of open source contributors, open source software stewards, and manufacturers. Under the PLD, the definition of 'product' has expanded to include software, potentially making manufacturers liable for defects in products that incorporate open source software components. The literature concludes that due diligence requirements now mandate a thorough assessment of open source software through documentation practices like security attestations, SBOMs, and model cards to ensure safety for commercial integration.

The Linux Foundation's Pathways to Cybersecurity Best Practices in Open Source report mentions that leading projects like CIP, Yocto, and Zephyr are already closing the gap to meeting CRA requirements through advanced security practices. The report defines the new role of open source software steward–organisations that support open source development without direct monetisation–and details their responsibilities for vulnerability handling and reporting. The literature recommends that projects establish Product Security Incident Response Teams (PSIRTs) and adopt semantic versioning to facilitate compliance for downstream manufacturers.

In light of this, The Black Duck 2026 OSSRA Software Governance in the AI Era report suggests that organisations now face urgent pressure to maintain current SBOMs and respond to exploitable vulnerabilities within 24 hours to meet the requirements of the EU Cyber Resilience Act (CRA).

The Open Source Security Foundation Blog Post on Preserving Open Source Sustainability argues that voluntary security attestation programs under the CRA must be designed to support manufacturers without burdening upstream volunteer developers. The author argues that proposals advocating for upstream self-attestation risk deviating from the 'no warranties, no liabilities' model that enables the scale of the open source ecosystem. Instead, effective compliance should prioritise automation and transparency, utilising build-time SBOMs and verifiable artifact provenance. The article emphasises that liability must rest downstream with the entities that commercially place products with digital elements on the market.

The OpenSSF Annual Report notes that to address the CRA, the foundation launched a Global Cyber Policy Working Group to provide a forum for aligning international cybersecurity regulations with open source practices. The Linux Foundation's Pathways to Cybersecurity Best Practices in Open Source notes that while leading projects (CIP, Yocto, Zephyr) have strong security practices, the CRA's mandatory five-year support period exceeds many projects' current commitments leaving a significant readiness gap.

4.2.2 AI is a dual-edged force

The Black Duck 2026 OSSRA Software Governance in the AI Era Report from March 2026, concludes that software development has entered a pivotal moment where AI-assisted coding is fundamentally altering the risk and compliance landscape. Open source software remains a universal foundation, present in 98% of audited codebases, yet mean vulnerabilities per codebase have grown to 581.

The report highlights an expanding attack surface where 65% of surveyed organisations experienced a software supply chain attack in 2025. Furthermore, a significant governance gap exists, as 76% of companies acknowledge that developers use AI assistants even against corporate policy, creating 'Shadow AI' risks.

4.2.3 Transparency is a New Mandate

There is broad consensus that transparency mechanisms - SBOMs, attestations, and verifiable provenance - are shifting from optional best practices to requirements. Sonatype's State of the Software Supply Chain 2026 report argues that software supply chains have reached 'machine scale,' with open source now serving as the internet's critical infrastructure and being reused more often. It frames trust at scale as the central engineering and business challenge of modern software. Key findings include open source malware operating as a nation-state business model that targets developer credentials, CI secrets, and build environments; vulnerability intelligence failing at the moment it matters most, with basic data often missing, late, or wrong; and avoidable consumption of known-vulnerable versions persisting even when fixes exist, driven by workflow inertia and unclear ownership. The report warns that AI-assisted development accelerates both productivity and risk, amplifying bad inputs at machine speed by selecting non-existent versions or unsafe packages without guardrails, and concludes that transparency through SBOMs, attestations, and provenance is now a mandate, with compliance shifting from evaluating policy documents to evaluating build outputs.

The OpenSSF SBOM Risk whitepaper from September 2025, pushes this further, suggesting that Software Bills of Materials (SBOMs) must be transformed from static documents into tools for active risk management. The guide identifies six core use cases for enriched SBOM data, including identifying End-of-Life (EOL) components and assessing licensing risks in both open and closed source software. Organizations are encouraged to enrich SBOMs with Vulnerability Exploitability eXchange (VEX) data to prioritise remediation efforts effectively. The report concludes that SBOM lifecycle management is becoming a shared responsibility that requires automated tools to facilitate data merging and continuous vulnerability monitoring.

4.2.4 Sustainability of open source and Maintenance

Underlying the security concerns is a structural fragility in how open source is maintained. The Linux Foundation's Census III of Free and Open Source Software: Application Libraries report notes that the health of the modern economy is deeply tied to open source application libraries, yet many critical projects lack a robust maintainer base. Among the top 50 non-npm projects, 17% have only one developer accounting for more than 80% of commits, creating significant systemic vulnerabilities. The report documents a 500% increase in the use of Rust components, signaling an industry-wide shift toward memory-safe languages. Furthermore, the report identifies a critical need for a standardised software component naming schema (such as pURLs) to improve supply chain transparency and vulnerability tracking.

The OpenSSF post on Sustainability Article and the Linux Foundation's 2025 ROI for Software Contribution report both grapple with this: regulatory burdens risk overwhelming volunteer developers.

The Linux Foundation report argues that contributing to open source provides a measurable strategic investment with a 2–5x return across engagement types. The research establishes that 45% of organisations maintain private forks, which costs them over 5,000 labor hours per release cycle and accumulates significant technical debt. Organisations that transition from passive consumption to upstream contribution see a 10% average increase in product development speed and gain at least two months of advance notice for critical project changes. Ultimately, the report argues that open source contribution is a core business driver that reduces Total Cost of Ownership (TCO) by distributing maintenance costs across the broader community.

4.2.5 From consumption to 'Stewardship'

The impetus is from consumption to stewardship. Throughout all reports, organisations are urged to shift from passive consumers to active stewards. The LF ROI report quantifies the cost of passivity (45% of organizations maintain private forks costing 5,000+ labor hours per release cycle) and the benefits of contribution (10% faster product development, advance notice of critical changes). Despite this move, the OpenSSF CRA Awareness and Readiness report found that 54% of respondents struggle to distinguish between the legal definitions of 'manufacturers' and 'stewards', which carry entirely different regulatory obligations

The OpenSSF Annual Report 2025 opines that the community has achieved collective progress through milestones such as the launch of the Open Source Project Security (OSPS) Baseline. In 2025, demonstrating collective investment in this direction, the foundation invested $5.8 million into 14 critical open source projects through the Alpha-Omega initiative to harden core infrastructure. The report also highlights advancements in AI model transparency and secure software delivery, including the release of trusted publishing for major package repositories, plus the OSPS Baseline and a Global Cyber Policy Working Group.

4.2.6 Converging on a single message - the ecosystem must invest

The literature converges on a single message - open source can only remain viable if the ecosystem collectively invests in security, transparency, and stewardship, while ensuring that new regulatory and security burdens are distributed fairly and do not crush the volunteer developers who make it all possible. Open source is critical but fragile infrastructure, present in 98% of codebases and underpinning the modern digital economy yet resting on a precarious maintainer base with structural weaknesses like single developers responsible for over 80% of commits in 17% of top non-npm projects. Meanwhile, the threat environment is worsening and automating, with much of the risk being avoidable but persisting through inertia and unclear ownership. At the same time, AI is accelerating both productivity and risk simultaneously, amplifying bad inputs at machine scale and introducing governance gaps like Shadow AI where developers bypass corporate policy.

Regulation (especially the CRA) is reshaping accountability. New legal roles - particularly the open source software steward and the expanded definition of product under the PLD are pushing liability and due-diligence obligations to the forefront, though readiness gaps remain. SBOMs, attestations, and verifiable provenance are now requirements, with compliance moving from evaluating policy documents to evaluating actual build outputs.

There is more uncertainty for open source developers and stewards. Investment must be made in supporting them, the OpenSSF CRA Awareness and Readiness report documents. While the CRA exempts non-commercial open source software, 61% of non-commercial developers remain unsure if the regulation affects them. 59% indicate that clear guidance on their CRA status is needed to reassure them about continuing to contribute. Furthermore, open source stewards identify a strong need for structural investments, specifically financial support (50%) and legal guidance (39%), to manage rapid vulnerability responses and build shared compliance infrastructure.

Ultimately, the literature shows that the path forward is shared responsibility and stewardship. Across all reports, the consistent prescription is automation, transparency, and a rebalancing of effort from passive consumption toward active, sustainable contribution. This is a shift that is both a security necessity and a measurable business driver.

5. UK Draft Cyber Regulation

The UK's Cyber Security and Resilience (Network and Information Systems) Bill was introduced to the Commons on 12 November 2025 following its announcement in the 2024 King's Speech. It is expected to receive Royal Assent later in 2026. The proposed UK Cyber Security and Resilience (Network and Information Systems) Bill builds upon the earlier UK NIS Regulation from 2018. It is the most significant update to UK cyber law since the 2018 NIS Regulations it replaces, and it represents the UK's first attempt to draft a cyber resilience regime in the post-CRA, post-NIS2 landscape.

The Bill is fundamentally different from product-focused regulation such as the EU's Cyber Resilience Act (CRA). The Bill will be Tabled before the House of Commons and receive its third Reading on 10th June.

5.1 What the Bill does

The 2018 Network and Information Systems Regulations (NIS) regime covered operators of essential services in sectors like energy, transport and healthcare. The bill extends the existing framework by adding medium and large managed service providers, data centres above one MW rated IT load (10MW for enterprise data centres, jointly regulated by DSIT and Ofcom), and designated critical suppliers whose failure would cascade into essential services.

In-scope entities must issue a light-touch notification within 24 hours of a significant incident and a full report within 72 hours. Regulators gain proactive investigation powers, cost-recovery mechanisms, and the ability to designate further entities by secondary legislation without returning to Parliament. Penalties scale to £17 million or four percent of global turnover for serious breaches, with daily fines of up to £100,000 for continuing non-compliance.

The Bill makes explicitly clear that a regulated organisation's security posture is inseparable from the security of what it consumes. Operators of essential services and digital service providers will be expected to manage cyber risk in their suppliers. In practice, this could mean SBOMs, vulnerability disclosure expectations, and provenance evidence flowing upstream from regulated buyers to everyone who sells to them.

5.2 What remains to be seen

The Bill makes no mention of open source. Provisions in respect to open source are likely to be implemented in secondary legislation at a later date.

Maintainers of software libraries used by regulated (and soon-to-be regulated) organisations have no direct duties or recognition within the letter of the law. That is not to say however that they aren't impacted by the bill, as maintainers and contributors could face indirect pressure through supply-chain obligations placed on the regulated entities which rely on their output.

Most of the operational aspects of the law will arrive through secondary legislation and regulatory guidance. The UK Government has said it intends to consult on implementation through 2026, with phased commencement following Royal Assent.

The Bill will come into force in phases once it has become an Act. Certain measures will come into force on Day 1 or on Month 2 following Royal Assent, while others will be brought into force through future secondary legislation (also known as "commencement regulations"), at a time determined by the Secretary of State.

UK Cyber Security and Resilience (Network and Information Systems) Bill Phases of Implementation

When	Provisions
Day One	• Future proofing • The post-implementation review
Month Two	• Statement of strategic priorities • Information sharing
Via secondary legislation	• Powers of Direction • Data Centres • Relevant Managed and Digital Service Provider updates • Large Load Controllers • Critical suppliers • Incident reporting • Costrecovery

Figure 4

6. Conclusion

Dr. Jennifer Barth
Chief Research Officer, OpenUK
Research Director, Symmetry & FSP

This report arrives at a pivotal moment for open source software, cybersecurity regulation, and digital infrastructure. Across the contributions, interviews, literature review, and regulatory analysis, one conclusion emerges clearly: open source is too important to remain invisible within cybersecurity policy, yet too complex to be regulated using models designed for physical products and traditional manufacturing. The challenge facing governments, regulators, businesses, and the open source community is how to incorporate open source without damaging the ecosystem on which modern digital society depends.

The findings demonstrate the scale of this challenge. Open source software now underpins almost every aspect of modern technology, from critical national infrastructure and healthcare systems to financial services, cloud platforms, and artificial intelligence. Yet among the twenty-three jurisdictions considered in the report, only two have meaningfully addressed open source within cybersecurity frameworks (EU and US). Most countries continue to regulate organisations rather than software itself, creating indirect pressures that flow upstream through procurement requirements, supply chain obligations, vulnerability management expectations, and software assurance requirements. As a result, open source is increasingly regulated in practice, even where it remains largely absent from legislation.

The European Union's Cyber Resilience Act (CRA) represents the most significant attempt yet to address this gap. Contributors throughout this report consistently identify the CRA as a landmark piece of legislation, not because it is perfect, but because it is the first major regulatory framework to recognise that open source occupies a distinct position within the software supply chain. The creation of the Open Source Software Steward category marks a historic shift. As both Mirko Boehm of Linux Foundation Europe and Rebecca Rumbul of the Rust Foundation explain, this recognition was not inevitable. It emerged only after extensive engagement by foundations, maintainers, policy organisations, and industry leaders who highlighted the potentially devastating consequences of treating volunteer contributors and non-profit foundations as commercial manufacturers.

Their experience offers an important lesson for future policymakers. Effective regulation requires an understanding of how software is actually developed. Open source is not produced through linear manufacturing processes. It is collaborative, iterative, and continuously evolving. Attempts to impose static concepts of production, certification, and liability onto dynamic software ecosystems risk creating obligations that are either impossible to meet or counterproductive to security outcomes. The eventual introduction of the steward model demonstrates that regulators can adapt when they engage directly with those building the systems being regulated. The challenge now lies in ensuring that the standards underpinning the CRA preserve this balance and do not unintentionally recreate the very problems that the stewardship model was designed to solve.

At the same time, the report highlights that regulation alone cannot solve the underlying security challenges facing open source. Throughout the contributions, a recurring theme is the fragility of the ecosystem itself. The literature shows that many critical components are maintained by small teams or even single individuals, despite underpinning services worth trillions of pounds globally. The Mythos Effect described by Andrew Martin has brought this fragility into sharp focus. AI-driven vulnerability discovery has compressed the time between vulnerability creation and exploitation to unprecedented levels, exposing weaknesses that were previously hidden by scale, complexity, or limited human capacity.

As Andrew Martin argues, AI has revealed the consequences of decades of underinvestment in security, maintenance, and resilience. Similarly, Adrian Mouat's emphasis on eliminating vulnerabilities through disciplined engineering practices and Liz Rice's focus on rapid mitigation and updates both point to a broader truth: the solution is not retreat from openness but investment in better operational practices.

This theme is reinforced by the discussion of Software Bills of Materials (SBOMs), provenance, and supply chain transparency. The evidence is clear that static compliance approaches are becoming obsolete. As AI-generated code becomes increasingly prevalent, organisations require continuous visibility into their software supply chains. SBOMs, attestations, vulnerability exploitability data, and build provenance are no longer optional best practices; they are rapidly becoming foundational requirements for managing risk. However, as several contributors note, transparency mechanisms alone cannot create security. They must be supported by organisational capability, automated remediation, secure development practices, and a culture of ongoing maintenance.

The report's examination of the NHS response to Mythos provides a practical illustration of these tensions. While contributors recognise the extraordinary responsibility carried by public sector organisations managing sensitive data, they broadly conclude that retreating from openness is not a sustainable security strategy. Clare Schramm argues that treating open source itself as the threat risks confusing visibility with vulnerability. This highlights a broader lesson: transparency and collaboration remain essential security tools, even in an era of increasingly capable AI systems.

Alongside regulation and security, the report identifies a third critical challenge: sustainability. Across the literature and contributor perspectives there is remarkable consensus that the open source ecosystem can no longer rely on informal goodwill alone. The economic value generated by open source has vastly outpaced investment in its maintenance. Andrew Martin's call to "pay open source maintainers yesterday" captures a growing recognition that organisations consuming open source at scale must contribute more actively to its health and security.

The emergence of AI agents introduces an additional dimension to this discussion. As explored in the chapter on identity and trust, organisations are entering a world where non-human actors can access systems, consume information, and make decisions at machine speed. The challenge is no longer simply authenticating users, but governing context, permissions, and authority.

Taken together, these themes point towards a common conclusion. The future of cybersecurity will depend on moving from passive consumption to active stewardship. Regulation, transparency, automation, AI governance, supply chain security, and maintainer sustainability are not separate challenges but interconnected aspects of the same problem. Organisations can no longer assume that security is someone else's responsibility.

7. Formalities

7.1 Authors

Professor Amanda Brock, CEO, OpenUK
OpenUK CEO, Amanda's built one of open source's most recognised and impactful organisations. Executive Producer of State of Open Con (2023- 2025), Amanda's a globally sought-after keynote speaker.A lawyer with 25 years' experience, 5 as GC of Canonical, she's been instrumental in shaping open source's legal frameworks, as she was internet law during the early 2000's. Regularly contributing to tech press, she edited 'Open Source: Law, Policy and Practice', (2022).

Recognition: Computer Weekly 50 Most Influential Women in UK Tech (2023, 2024); Computing IT Leaders 100 (2023, 2024); Lifetime Achievement Award WIPL (2022); Women Who Will Changemaker (2023); INvolve Heroes (2022, 2023); Novi Awards (2024) and Ambassador, Open Charge Alliance.

Advisory Appointments: UK Cabinet Office Open Standards Board; UKRI Digital Research Infrastructure; UKRI Exascale; KDE; commercial boards – Mimoto, Scarf, FerretDB and Space Aye; and is Fellow Open Forum Academy; Distinguished Fellow Rust Foundation; and European Representative, OIN.

Dr. Jennifer Barth, Head of Research, Symmetry Research at FSP and Chief Research Officer, OpenUK
Jenn holds a DPhil from the University of Oxford and is Head of Research at Symmetry Research at FSP, where she brings a rare blend of ethnographic depth, commercial acuity, and ethical clarity to the AI conversation. Jenn's work sits at the intersection of technology, organisational culture, and socioeconomic change. She is known for translating complex technological shifts into human-centred insights that resonate deeply with executive audiences. Her measured, critically informed perspective — grounded in qualitative research and lived experience — enables leaders to move from abstract concern to practical, principled action. Alongside her role at FSP, Jenn serves as Chief Research Officer at OpenUK.

Karan Saini, Research Manager, OpenHQ
Karan Saini is a New Delhi-based security researcher, public interest technologist, and open source contributor. He was most recently a Senior Information Controls Fellow with the Open Technology Fund, hosted by the Internet Governance Project at the Georgia Institute of Technology, where he researched the scale of DNS-based web censorship in India. He regularly writes for leading publications and has contributed to award-winning media projects, such as Eco-Bot.Net, a net-art project active during the COP26 climate talks (British Journalism Awards - Innovation of the Year 2022) and the BBC documentary "The Trap" (nominated for the 2024 International Emmy Award in Current Affairs.)

7.2 Contributors

Adrian Mouat, DevRel Engineer, Chainguard
Adrian Mouat is a DevRel Engineer at Chainguard and the author of the O'Reilly book Using Docker. He brings deep expertise in containers, Kubernetes, registries, and supply chain security. Adrian is also the creator of the open source Trow registry, a frequent conference speaker, and a passionate contributor to DevOps and cloud native communities.

Andrew Martin, CEO, ControlPlane
Andrew is Co-Founder and CEO, ControlPlane, and has an incisive security engineering ethos gained building and destroying high traffic web applications. Proficient in systems development, testing, and operations, he is at his happiest profiling and securing every tier of a cloud native system, and has battle hardened experience delivering containerised solutions to enterprise and government.

Chaamini Mangaleswaran, Lead Sales Engineer, WSO2
Chaamini is a Lead Sales Engineer at WSO2 with over a decade of experience in engineering and solution architecture. She works closely with enterprise customers across multiple industry verticals, helping them architect modern cloud-native applications, navigate complex architectural challenges, and accelerate their digital transformation journeys.

Clare Schramm, MD, Technology Platforms, Lloyds
Clare Schramm Fergus is a technology platforms MD at Lloyds Banking Group. She runs a team of engineers responsible for the design, build, and evolution of core banking systems using AI. Previously, Clare was the Cloud Engineering MD at BT where she ran the technology platforms that hosted BT's networks. She also worked for two hyperscalers, Google and IBM, in various engineering leadership roles. Clare is a systems engineer by background.

Liz Rice, Chief Open Source Officer, Isovalent
Liz Rice is Chief Open Source Officer with eBPF specialists Isovalent, creators of the Cilium cloud native networking, security and observability project. She was Chair of the CNCF's Technical Oversight Committee in 2019-2022, and Co-Chair of KubeCon + CloudNativeCon in 2018. She is also the author of Container Security, published by O'Reilly.

Matt Barker, CEO, BoltMCP
Matt Barker is co-founder and CEO of BoltMCP, focused on helping enterprises securely connect AI to their data and control how it's used. Previously, he was CEO and co-founder of Jetstack, where he helped build one of the most widely adopted projects in the cloud native ecosystem, cert-manager. Under his leadership, Jetstack bootstrapped to a team of 50, and was acquired by Venafi in 2020.

Matt Jarvis, Director of Applied AI, Snyk
Matt Jarvis is a Director of Applied AI at Snyk. Matt has spent more than 20 years building products and services around open source software, on everything from embedded devices to large scale distributed systems.

Mirko Boehm, Senior Director, Community Development, Linux Foundation Europe
Mirko Boehm contributes to free and open source software as a community builder, licensing expert and researcher. His work spans major projects including the KDE Desktop, the Open Invention Network and the Open Source Initiative. He holds a PhD in innovation economics and serves as visiting lecturer and researcher on open source software at the Technical University of Berlin.

Rebecca Rumbul, CEO, Rust Foundation
Rebecca is the Executive Director and CEO of the Rust Foundation, a global non-profit stewarding the Rust language, supporting maintainers, and ensuring that Rust is safe, secure, and sustainable for the future. She holds a PhD in Politics and Governance, and has worked as a consultant and researcher with governments, parliaments and development agencies all over the world, advocating for openness and transparency, and developing tools to improve digital participation.

Sal Kimmich, CEO, GadflyAI
Sal Kimmich is CEO and cofounder of GadflyAI. Sal's work is centered on advancing secure computing practices and pioneering privacy-enhancing technologies to safeguard sensitive data. At the Confidential Computing Consortium (CCC) within the Linux Foundation, they play a pivotal role in the development and implementation of Open Source projects that establish Trusted Execution Environments (TEEs)

7.3 About the Creators of this Report

OpenUK

OpenUK is the unique open tech industry organisation for the business of open technology in the UK. It spans the opens – software, hardware, data, standards and AI and is the convening point for the UK's business, academic and contributing communities across open tech. Our work supports the UK's journey to become "The State of Open". Our organisation is run with the support of our volunteer community and their leadership in the tradition of open source manner delivering on three pillars: community, legal and policy and learning. OpenUK's Community is recognised through our world-leading recognition programme including the Open UK Awards (the Oscars of Open Source) now in their 7th year, New Year's Honours Lists and Ambassador Scheme.

OpenUK undertakes research and reporting both on its own account through its "State of Open Reports" and on a commissioned basis for third parties. Case studies, Thought Leadership, Surveys and desk-based research are included in our reporting which pushes the envelope and leads the way. Our Research and Reporting Show and Tell events coalesce the global open source research communities digitally to regularly update and share research practices and topics.

The community's strength is channelled to enable a cohesive voice that responds to legislative proposals and sets policy. We have set the agenda in policy matters across openness in the UK and beyond. OpenUK's Policy work leads the conversations around open source licensing and commercialisation, AI openness and cloud computing and other key topics in open source, as they emerge. Engagement with UK policy makers is supported by a volunteer Policy Advisory Board and by experts in our volunteer Advisory Boards and the open source communities. Our Advisory Boards span AI, Communications Tech, Data, Finance, Hardware, Healthcare, Security, Software, Space, Sustainability and Quantum Computing. We are able to provide in-

dustry experts from the commons for speaking engagements, consultancy and advisory boards.

OpenUK is the second organisation established anywhere in the world with open source policy as its purpose, our approach is holistic to and representative of the entire open ecosystem. OpenUK undertakes a broad range of activities in support of its policy work and is a day one member of GaiaX and UK's GaiaX Hub Coordinator, hosted one of the biggest tech events at COP26, and was the first organisation in open tech to put a Sustainability Policy and Chief Sustainability Officer in place. Skills and Learning form our third pillar and our Learning work has spanned initiatives for children including our award winning Kids Camps which teach coding, open source and sustainability in a real world context; and exploring the business of open source through our Founder training. We have shared several hundred hours of digital training. Our ambitions include a UK apprenticeship module and adding open source to the UK curriculum.

The State of Open Con has become one of the world's leading open source conferences since its inception by OpenUK in 2023. In 2027 we expect to host 1000 people throughout 8 tracks and plenary sessions, with at least 50 partners in our delegate experience space and over 200 speakers in London. For 2026 we are on the road, with smaller events across the UK, meeting our audience. The first of these took place in Edinburgh on 5 June 2026.

Our small events team deliver to the highest standards a series of unique events throughout the year and our community organise UK-wide OpenUK meet-ups. Contact OpenUK mailto:admin@openuk.uk

Symmetry Research

Symmetry Research, an FSP company, looks beyond the surface and behind the curtain of the fundamental innovations and trends shaping our society, markets, culture, and values. We are academics and researchers looking at the intersections of emerging technology and socioeconomic impact, producing independent research for thought leadership and business solutions. We take a research-led approach to helping organisations transform through and with technology.

Symmetry Research's mission is to share and grow knowledge about the interaction of technology and everyday lives. We want to understand the past, present, and future of human interaction with emerging technologies and socioeconomic changes-from behaviour to context, nature to nurture, origin to experiences-we do independent research for high-stakes business decisions.

7.4 Methodology

The research used a mixed method approach to explore and demonstrate the state of policy and regulations in relation to cyber security globally. Interviews were conducted with industry leaders, founders and open source software experts and included as case studies and thought leadership. Occasio has provided us with input on our State of Open Con event on 8th June.

Occasio

Occasio provides the digital infrastructure for organisations to capture, reference, and exchange short-form textual insights seamlessly. Our platform transforms fragmented observations into a high-impact content library, using AI-powered synthesis and integrated citation management to turn raw data into powerful narratives. From strategic soundbites to key selling points, Occasio ensures your collective intelligence is centralised, standardised, and ready for delivery. By bridging the gap between scattered internal data and actionable intelligence, we help teams move beyond the noise to build a lasting legacy of knowledge.

7.5 References

- Black Duck Software, Inc. "2026 Open Source Security and Risk Analysis Report: Software Governance in the AI Era." March 2026. https://www.blackduck.com/content/dam/black-duck/en-us/reports/rep-ossra.pdf
- Boehm, Mirko, Hilary Carter, and Cailean Osborne. "Pathways to Cybersecurity Best Practices in Open Source: How the Civil Infrastructure Platform, Yocto Project, and Zephyr Project are Closing the Gap to Meeting the Requirements of the Cyber Resilience Act." The Linux Foundation, March 2025.
- Colonna, Liane. "The End of Open Source? Regulating Open Source Under the Cyber Resilience Act and the New Product Liability Directive." Computer Law & Security Review 56 (2025): 106105. https://doi.org/10.1016/j.clsr.2024.106105
- European Commission. Cyber Resilience Act. Accessed April 21, 2026. https://www.cyberresilienceact.eu
- European Commission. "Cyber Resilience Act: Proposal for a Regulation on Cybersecurity Requirements for Products with Digital Elements." Brussels, 2022.

- Lawson, Adrienn, and Sam Boysel. "ROI for Open Source Software Contribution: Insight from the Open Source ROI Survey and Economic Model." The Linux Foundation, February 2026. https://www.linuxfoundation.org/research/contribution-roi
- Linux Foundation. https://www.linuxfoundation.org/research/cra-compliance-best-practices
- National Institute of Standards and Technology (NIST). Secure Software Development Framework (SSDF). Gaithersburg, MD: NIST, 2022. https://csrc.nist.gov/projects/ssdf
- Nocera, Sabato, Simone Romano, Massimiliano Di Penta, Rita Francese, and Giuseppe Scanniello. 2025. "On the Adoption of Software Bill of Materials in Open-Source Software Projects." Journal of Systems and Software 230: 112540. https://doi.org/10.1016/j.jss.2025.112540
- Nagle, Frank, Kate Powell, Richie Zitomer, and David A. Wheeler. "Census III of Free and Open Source Software: Application Libraries." The Linux Foundation, December 2024. https://www.linuxfoundation.org/hubfs/LF%20Research/lfr_censusiii_120424a.pdf?hsLang=en
- Neag, Madalin. "Preserving Open Source Sustainability While Advancing Cybersecurity Compliance: Reflections on Voluntary Attestation Models Under the Cyber Resilience Act." Open Source Security Foundation (OpenSSF), January 2026. https://openssf.org/blog/2026/01/21/preserving-open-source-sustainability-while-advancing-cybersecurity-compliance/
- Open Source Security Foundation (OpenSSF). "2025 Annual Report." December 2025. https://openssf.org/wp-content/uploads/2026/01/2025_Annual_Report_0130.pdf
- Open Source Security Foundation (OpenSSF). "Improving Risk Management Decisions with SBOM Data." September 2025. https://openssf.org/resources/improving-risk-management-decisions-with-sbom-data/
- Open Source Security Foundation (OpenSSF) "2026 CRA Awareness and Readiness." June 2026. Linux Foundation (LF). "Unaware and Uncertain: The Stark Realities of Cyber Resilience Act Readiness in Open Source." March 2025. https://www.linuxfoundation.org/research/cra-readiness
- Open Source Security Foundation (OpenSSF). "SBOMs in the Era of the Cyber Resilience Act: Toward a Unified and Actionable Framework." 2025. https://openssf.org/blog/2025/10/22/sboms-in-the-era-of-the-cra-toward-a-unified-and-actionable-framework/
- OWASP Foundation. "CycloneDX Specification" Accessed April 21, 2026. https://owasp.org/www-project-cyclonedx
- SOCRadar. "UK Threat Landscape Report." Accessed April 21, 2026. https://socradar.io/wp-content/uploads/2025/06/UK-Threat-Landscape-Report-2025.pdf
- UK Government. Cyber Security and Resilience Bill. Accessed April 21, 2026. https://www.gov.uk/government/collections/cyber-security-and-resilience-bill

7.6 Sponsors

This report does not have a single sponsor, which OpenUK and OpenHQ consider important for general industry reporting of this nature, securing our independence.

Our research team is available to create sponsored reports and regularly hosts small group dinners to support short papers, which are sponsored. However, our independence in reporting is critical to us and our sponsors never set the agenda beyond participating in the topic discussion.

If you would like to find out more contact mailto:admin@openuk.uk

We are grateful to our general sponsors who make our work possible including our industry reporting

Reports: https://openuk.uk/report/

Newsletter: https://openuk.us4.list-manage.com/subscribe?u=9d6308e45152bb-731122903de&id=34b8a50cce

State of Open Con Edinburgh: https://stateofopencon.com/edinburgh-soocon26/