



Research



Europe

Sponsored by



World of
Open Source

2026 State of Open Source in Europe

Navigating AI and Digital Sovereignty
Through Open Governance

*Dr. Mirko Boehm, Linux Foundation Europe
Paula Grzegorzewska, Linux Foundation Europe
Marco Gerosa, PhD, Northern Arizona University*

Foreword by
Greg Kroah-Hartman, *The Linux Foundation*

October 2026

2026 State of Open Source in Europe

94% of European organisations think that digital sovereignty is strategically important to their organisation, with 61% calling it very important.



Security and privacy lead as the highest driver of sovereignty, at 71%.

Advancing national and regional interests through sovereignty is felt most acutely in Europe, where its importance is 12 percentage points higher than the rest of the world.



The Tech Sovereignty Package is the first major global legislative bet on open source as a pillar of strategic autonomy, backed by a planned €2bn in funding over seven years.

Open source returns 4.1x benefit-to-cost on average, rising from 3.6x with no formal governance to 4.6x at the most mature governance level.



59% of European organisations contribute code upstream, but only 37% participate in open source project governance.



Foundation-hosted, neutrally governed projects are Europe's preferred model at 52%, 8 percentage points ahead of the rest of the world.



83% of contributors get advance warning before a breaking change lands, one of several concrete paybacks participation affords beyond code.



94% of European organisations are already using or piloting generative AI coding tools.



48% of European organisations report more open source use because of AI coding assistants, against just 5% reporting less.



76% of European organisations say AI tools help their team get more value from the open source they already run.



Only 28% of European organisations expect AI coding tools to reduce their reliance on open source in future.

Contents

Foreword	4
Executive summary	5
Introduction: Open source trends among European organisations	6
The impact of AI on the open source ecosystem	7
Open source and AI sovereignty	8
AI's long-term effect on open source adoption	10
AI and cybersecurity	13
Community and technical governance	15
Governance maturity	16
Does contributing upstream actually buy influence?	20
The real cost of maintaining private forks	22
What actually blocks open source adoption	23
Digital sovereignty and technological resilience	25
EU regulatory frameworks: the Tech Sovereignty Package	28
Public sector and procurement	30
Other policy developments in 2026	31
Digital policy moves towards verticals	33
Conclusions	34
Methodology	35
Demographics	35
About the authors	37
Acknowledgments	38

Foreword

As a Linux kernel developer, maintainer, and member of the kernel security team, I've seen a wave of new contributors to our project over the past nine months. The majority of these contributions are fixes based on reports from various LLM tools. These tools are churning out reports in which they claim hundreds of bugs in codebases. Almost always these inflated numbers turn out to be incorrect once reviewed by developers who know the code. This has presented a major challenge to our maintainers and reviewers, and Linux is not alone in facing this onslaught.

These submissions can be helpful, many times they are fixing real bugs that have persisted for many years in established codebases. The rise in Linux kernel CVEs issued this year, from an average of 55 per week before the tools started, to the current rate of 230 per week, shows their value in making open-source projects more secure and bug-free. But these submissions also create a lot of extra work for reviewers and maintainers, because the tools often produce incorrect results that can easily fool developers, thanks to the “wall of text” LLMs generate by default.

Over the summer I worked with a number of graduate students in the Computer Security program at Vrije Universiteit Amsterdam to test just how valid these tools' output is. We saw that even for the best and most powerful LLM tools available, almost half of the changes that were produced were incorrect. Because of this, I'd recommend anyone producing these types of changes be extra careful in reviewing and testing before submitting them, so as not to waste a reviewer's time.

Because these submissions are so often wrong, Linux maintainers are pushing back, requiring submitters to prove a change is correct before it's accepted. I strongly suggest that other projects do the same. The burden of validation should sit with the submitter, not the receiver, who is already carrying an increased workload because of this wave of submissions. Extra diligence on the submitters part is needed.

The open source ecosystem has faced challenges like this before, with the rise of static analysis tools and the spread of fuzzers, and the same will happen here with the wave of LLM-generated fixes. Eventually we will get through them all, and our codebases will be stronger for doing so. In the meantime, it's a long slog, but at least we are all in it together.

By Greg Kroah-Hartman, *Fellow, Linux Foundation; EU Prize for Excellence in Open Source, 2026*

Executive summary

This is the fifth edition of the Linux Foundation's annual State of Open Source in Europe research, combining survey responses from 367 European organisations with interviews conducted with open source maintainers, policymakers and practitioners across the continent. 2026 is the year digital sovereignty and AI's disruption of the open source ecosystem stopped being talking points and became operational mandates. Both are now being tested against a governance capacity that many European organisations still lack.

Artificial intelligence's effect on the ecosystem is already large and cuts both ways. Ninety-four percent of respondents already use or are evaluating generative AI, 48 percent report more use of open source libraries and frameworks because of it, and 76 percent say AI tools help their team get more value out of the open source they already run. Two worries complicate that picture: a bifurcation risk, in which large, well-resourced projects benefit while the long tail hollows out, and a governance lag, in which Europe writes a large share of the code behind AI systems without a seat at the table where the standards for those agents get set.

Governance maturity is the report's throughline. Thirty-one percent of European organisations report no formal open source governance at all. Organisations with no governance report a benefit-to-cost multiple of 3.6 times from their open source use, those with advanced governance report 4.6 times, against a European mean of 4.1 times, and upstream contribution itself returns between 2.3 and 3.3 times. Not capturing that higher return is not a neutral gap, it is a competitive one: an organisation extracting 3.6 times from the same technology a peer extracts 4.6 times from operates at a real, compounding disadvantage. Forty-eight percent of organisations maintain private forks, and

cannot say what that costs them. Adoption blockers cluster tightly around security (39 percent), licensing (36 percent), support (35 percent) and leadership buy-in (34 percent). This indicates the gap is procedural, not technical.

Digital sovereignty has undergone the same shift from talking point to mandate. Ninety-four percent of European organisations call it important to their technology strategy, 61 percent very important, with security and privacy (71 percent) and vendor lock-in and cost risk (55 percent) leading as drivers. The EU's Tech Sovereignty Package and Cyber Resilience Act are the first major policy bets on open source at this scale, arriving as Europe relies on non-EU providers for over 80 percent of its digital products, services, infrastructure and IP. The money is real but modest: the Open Source Strategy commits €2 billion over seven years, against the European public sector's own €264 billion in annual proprietary IT spending, a sum that becomes a strategy rather than a rounding error only if functioning OSPOs and enforced procurement rules show up alongside it.

Read together, AI, governance and digital sovereignty are not three separate stories but three views of the same capacity gap: organisations that build the governance capacity to engage with open source deliberately capture more of what AI, contribution and procurement all have on offer, while the rest are left responding to decisions made elsewhere. Europe's long-term resilience, on the evidence gathered here, depends not on building isolated regional silos but on deepening its participation in the open source commons it already helped build. Whether it does so from a seat at the table, or from outside looking in, is what the next twelve to twenty-four months will decide.

Introduction: Open source trends among European organisations

This report is the 5th edition of our annual research on the state of open source in Europe. Across these annual editions, Linux Foundation Research tracks Europe's transition from widespread software adoption to policy-backed digital sovereignty. The 2022 and 2023 reports established baseline usage metrics, highlighting community momentum alongside persistent gaps in formal Open Source Program Office governance. The 2024 report broadened the scope to organisational maturity, detailing how institutions navigate emerging open AI standards and regulatory rules under the EU Cyber Resilience Act. By 2025, the study showed the transition of open source into a part of critical digital infrastructure in Europe. This year, we continue expanding on this topic, but add considerations on the impact of AI and the importance of open governance in the context of digital resilience and global collaboration.

Nearly a third of European organisations report no formal open source governance at all. Governance maturity determines the scale of problem an open source community, or the organisation that depends on it, can actually solve, and that is the thread running under the three chapters that follow: whether community collaboration standards hold up, whether human judgment stays in charge as the tools improve, who protects maintainers who were already overburdened before AI arrived. We discuss this against the avalanche of AI innovation and the arrival of the Cyber Resilience Act which look like the seismic shifts of the moment, forces poised to remake the open source ecosystem across Europe and beyond. However, what stands out is not just the disruption itself but our ecosystem's remarkable capacity to adapt and evolve as the currents around it change. This report is offered in that

spirit – as guidance in a period of uncertainty.

The first thread is AI's effect on the open source ecosystem. It cuts both ways. Coding assistants lower the barrier to contribute and appear to be lifting open source adoption overall, yet the same tools generate more vulnerability reports than maintainers can triage, and leave Europe writing a large share of the code behind AI agents without a seat at the table where the standards for those agents get decided. This chapter follows that tension through open source's uneven progress toward AI sovereignty, AI's longer-run effect on how organisations actually adopt and rely on open source, and the AI-specific edge of the cybersecurity question, where AI accelerates the discovery of vulnerabilities and the noise around them at the same time.

The second thread is community and technical governance. Governance maturity, upstream influence, private forking and adoption blockers turn out to be not four separate problems but four views of one underlying capacity gap, concentrated specifically in EU-headquartered public sector organisations rather than the European sample at large. The chapter asks how mature European engagement in open source governance actually is, whether contributing upstream buys real influence over the projects an organisation depends on, what maintaining private forks instead actually costs, and what blocks adoption in the first place when the technology itself is rarely the problem.

The third thread is digital sovereignty and technological resilience, which has moved this year from a policy talking point to an operational mandate, chiefly through the EU's Tech Sovereignty

Package. The chapter traces that shift through the Cloud and AI Development Act and the Open Source Strategy, public procurement as digital sovereignty's most promising demand-side lever, the wider 2026 policy calendar including the Cyber Resilience Act's compliance clock, sovereignty's uneven reach into vertical industries such as energy and open, and RISC-V-based silicon under Chips Act 2.0. However, sovereignty is not a synonym for isolation. Europe's resilience, on this chapter's evidence, depends on deepening its participation in global open-source ecosystems, not retreating from them.

A few threads recur across all three chapters. Digital sovereignty has succeeded at the infrastructure layer, control over cloud, containers and data locality, while lagging markedly at the AI layer. Governance gaps and AI-adoption gaps concentrate in the same EU-headquartered public sector organisations. And the Cyber Resilience Act's compliance clock, like AI's effect on maintainer workload, acts as a forcing function for an overall better cybersecurity posture.


The impact of AI on the open source ecosystem

Artificial intelligence's adoption has been explosive, and its impact on the open source ecosystem is already disruptive and multi-layered. It reaches into the discipline of software engineering itself, into how people collaborate both on writing new code and on maintaining what already exists, and even into how the community thinks about intellectual property: the logic of copyright was built to create exclusivity because creativity was scarce, a premise that erodes as AI makes producing software drastically cheaper. The effects cut both ways. AI makes contributing easier, yet maintainers report feeling more overwhelmed than before. It helps root out bugs and surface security issues at a scale no human team could match alone, yet that same scale makes it harder for developers to separate the signal from the noise.

The survey data behind this report reads, at first glance, as good news. 94 percent of respondents already use or are evaluating generative AI, 48 percent report more use of open source libraries

FIGURE 1

AI adoption and its impact on open source



Source: 2026 World of Open Source Survey, Q25 (n = 236), Q26 (n = 220), and the displayed Q27 (n = 222), Europe only, excludes unsure.

and frameworks because of it, and 76 percent say AI tools help their team get more value out of the open source they already run (Figure 1). But a simple story, “AI is good for open source,” does not survive contact with the interviews conducted for this report. Two worries paint a more critical picture: a bifurcation risk, in which large, well-resourced projects benefit while the long tail of smaller ones hollows out, and a governance lag, in which Europe writes a great deal of the code but is not yet in the room where the standards for AI agents, governance, and collaboration get decided.

Where there is light, there is also shadow. Two questions recur throughout this chapter. Can the balance between vulnerabilities discovered and vulnerabilities fixed hold as AI accelerates both sides of that ledger at once? And can maintainers, already stretched before AI arrived, adapt to the new environment and the resulting workload it generates?

Open source and AI sovereignty

For decades, the animating project of open source software was to give users control over how, where, and by whom their software processed their data. That project appears suspended in the age of AI. Users and organisations alike now transfer undefined quantities of personal and proprietary data to AI models hosted in someone else’s cloud, often without asking who else can see it or what happens to it afterwards. Why take that risk?

Interviewees offer several familiar explanations. Excitement about a new technology outruns risk awareness, just as it did in the early years of social media. Regulation lags behind invention: rules are hard to write for a technology that is still evolving at breakneck speed. And a pragmatic preference for open weights over open source licences with strict four-freedoms guarantees has quietly displaced a more principled stance. Early adopters describe their own choice candidly. “I know this is immature tech, but I want to

learn it,” one interviewee said; another reasoned that giving AI a strict harness makes the risk manageable; a third argued that not learning and evolving with AI today only makes it harder to catch up tomorrow.

There are good reasons for caution. AI, at bottom, is excellent pattern matching. It is a powerful tool, but the burden of interpreting its output, hallucinations included, still falls on the human operator. This requires experience that often only senior staff can provide. Few consumers are currently able or willing to run even open source or open-weight models locally. There is evidence that models leak users’ data during ordinary use back into their own training, effectively absorbing it into the model’s embedded knowledge. And the ability to switch providers or models, once an organisation has committed to a particular AI stack, remains very limited.

Open source’s most fundamental effect on the technology sector was to hand the means of production to anyone with a computer and the willingness to learn. At least at this moment, this is not the case for the development of AI tools and models. On the one hand, it lowers the barrier to contribute, through coding assistants, better investigative tooling, and knowledge transfer mediated by conversational interfaces. On the other hand, building the means of production itself, the large language model that sits behind every assistant, now requires capital expenditure on hardware, energy and highly trained people that a laptop and a weekend cannot supply. While it is possible to use open source or open weights models, creating them is prohibitively expensive for individual contributors and decentralised communities. The low-resource collaboration that made open source successful in the first place does not transfer easily into the AI era.

Europe experiences this as a race it is not sure it is running. It is worth questioning the premise: much of the public discussion assumes that whichever party invents artificial general intelligence

(AGI) first will dominate forever, an assumption more common outside technical circles than within them. Whether or not AGI is around the corner, open and open-weight models currently trail the leading proprietary ones by roughly six months, not by years. It is precisely open source's competition-enabling role, its history of preventing single incumbents from locking down a market, that offers a real, if only medium-to-long-term, chance to stop the flipping markets pattern repeating in AI.

The survey confirms this concern about monopolistic control: 94 percent of European organisations call digital sovereignty important to their technology strategy, 61 percent very important (see Figure 21 in the digital sovereignty section). Public sector actors may have a stronger sensibility towards data protection and privacy (Figure 2).

Without powerful AI models created by an open community or domestic alternatives, AI sovereignty specifically lags behind infrastructure sovereignty, where open source has already solved the equivalent problem for operating systems, web servers and databases. Even Brussels appears to have made its peace with the situation. The EU's Cloud and AI Development Act speaks of "open source frontier models" as a policy target, not of a European-built proprietary one. Read charitably, that is pragmatism: a model six months behind the frontier is close enough for almost every real use case. Read less charitably, it concedes that Europe will not be the one training frontier models, only adapting and deploying what others release.


Either way, as the underlying models commoditise, the fight over openness does not disappear. It simply moves up a layer, to the harness: the agents, orchestration frameworks and protocols that decide how models actually get used. The Linux Foundation's new Agentic AI Foundation, launched on 9 December 2025 around Anthropic's Model Context Protocol, Block's goose and OpenAI's AGENTS.md specification, is that harness layer taking institutional shape in real time. Every founding Platinum member is American: Amazon Web Services, Anthropic, Block, Bloomberg, Cloudflare, Google, Microsoft and OpenAI. The only three European members, SUSE, SAP, and Ericsson, sit at lower tiers, Gold and Silver. It is the same pattern Europe already lives with in the Cloud Native Computing Foundation: strong on code contribution, thin on governance seats. This time, though, the pattern is visible from the foundation's first governance meeting.

Akrites, a second Linux Foundation initiative, launched on 25 June 2026 and seeded by the Alpha-Omega fund, coordinates confidential remediation of vulnerabilities in critical open source projects before they can be exploited. It was established as a direct response to frontier AI models that can now scan a major project and surface exploitable flaws in minutes, collapsing the expertise gap that used to separate attackers from defenders. Its

FIGURE 2

Public-sector employees emphasise national and regional sovereignty interests

National or regional interests as a driver of digital sovereignty



Source: 2026 World of Open Source Survey, Q19 by Q3, Europe only, percentage of European respondents citing national or regional interests as a driver of digital sovereignty, n = 229, DKNS responses and respondents with no Q19 selection are excluded. Occupation groups with fewer than 10 valid responses are omitted.

founding members include Ericsson and Vodafone alongside the same American cloud and AI companies that anchor the Agentic AI Foundation, since Akrites has no tiered membership structure.

Regulation offers a partial safety net here. General-purpose open source AI models are exempt from most obligations under the EU AI Act, but not from all of them: copyright policy, training-data summaries and systemic-risk duties still apply once a model crosses a capability threshold. For high-risk systems, the exemption disappears entirely: no open source carve-out at all, full conformity assessment and CE marking required from 2 August 2026. The practical effect is that the exemption is most generous exactly where the sovereignty stakes are lowest, and vanishes exactly where they are highest. A genuine policy response is nonetheless taking shape. The EU Open Source Strategy, announced on 3 June 2026 as part of the Tech Sovereignty Package, commits €2 billion over seven years, stands up an EU OSPO Network and a Digital Commons European Digital Infrastructure Consortium (EDIC), and directs €500 million into open RISC-V hardware. The Cloud and AI Development Act pairs with it on the procurement side, encouraging an open-source-software-first approach for public cloud and AI purchasing. It is worth resisting the temptation to call the problem solved: €2 billion over seven years set against the EU's own €264 billion in annual spending on proprietary IT looks less like a strategy than a rounding error, unless the institutional follow-through, functioning OSPOs, enforced procurement rules, and an EDIC that actually convenes, show up alongside the money.

Not every interviewee frames this as “AI sovereignty” specifically. Some describe it simply as vendor lock-in, with open source functioning as the switching mechanism that keeps an organisation from depending on any single infrastructure provider. That framing deserves a place alongside the more AI-specific arguments above:

infrastructure independence and AI independence are, in the end, versions of the same underlying problem.


AI's long-term effect on open source adoption

Coding assistants and other AI tools seem to tilt open source adoption upward: 48 percent of survey respondents report slightly or significantly more use of open source because of them, 47 percent report no change, and only 5 percent report using less (Figure 3). The long-term effect looks very likely net positive, though AI has not yet been transformative for most organisations.

What survey participants highlighted specifically reinforces that

FIGURE 3

Coding assistants are driving greater open source adoption



Source: 2026 World of Open Source Survey, Q26. Europe n = 220 (excl. 23 “unsure”); Rest of World n = 320 (excl. 33)

picture: 76 percent agree that AI tools help their team get more value out of the open source software they already use (Figure 4). Interviewees strike a similarly cautious note of optimism, a net positive effect overall, tempered by real concerns over licensing, AI-slop and maintainer burden that recur throughout the rest of this section

FIGURE 4

AI is strengthening organisations' engagement with open source

To what extent do you agree or disagree with the following statements about AI and open source?



Source: 2026 World of Open Source Survey, Q27. Europe only, n = 187-222 (excl. 21-56 "N/A or unsure", by row). Each donut is the share who agrees strongly or somewhat with that statement

The reasons for that optimism are fairly consistent, though not everyone shares it: only 28 percent agree that AI coding tools will reduce their organisation's overall reliance on open source in future (Figure 4). One mechanism is a reduced dependence on scarce specialist expertise: coding assistants lower the bar enough that more contributors can meaningfully work on more

projects, a form of democratisation. A second is amplification (rather than substitution) for experienced technologists, who already understand how the systems they build should behave and use AI to fill in the details rather than to replace their judgement. A third is a genuinely lower barrier to entry for new or occasional contributors: new contributors to the Linux kernel tripled in the second quarter of 2026¹. Sixty percent of respondents report that AI tooling has made it easier for their organisation to contribute to open source projects at all, whether through documentation or pull request review. Maintenance is a different, less settled story. Coding assistants do not get tired or bored, which suits them to the routine work humans tend to dislike: documentation, test generation, issue triage. Yet only 55 percent of respondents report having used AI to accelerate open source maintenance this way. Maintenance automation is real, but it remains uneven and considerably less mature than usage or contribution.

Risks remain. Junior roles are under pressure, and with them the pipeline that is supposed to produce the senior developers of the next decade: Europe's entry-level technical roles contracted by 3 percent in 2025 while the rest of the world expanded junior hiring by 14 percent, and the mid-senior talent European organisations will need by 2036 is not being trained today². At the same time, open source engagement reportedly improves skills of engineering teams, as claimed by 76% of respondents of our survey, while 53% state that hiring and retention have become easier (Figure 5). Technical teams have long recognized that participating in open collaborative workflows develops technical expertise, which will potentially become even more vital in the age of AI. Generally, greater OSS engagement is linked to stronger talent attraction.³


¹Development statistics for the 7.2 kernel, <https://lwn.net/Articles/1088776/>

²Mirko Boehm, "Who will be the senior engineers of 2036?", The Linux Foundation, 24 June 2026, <https://www.linuxfoundation.org/blog/who-will-be-the-senior-engineers-of-2036>

³Marco Gerosa and Adrienn Lawson, "The State of Global Open Source 2025", foreword by Jon Seager, The Linux Foundation, October 2025

FIGURE 5

Open source strengthens hiring and engineering capability



Source: 2026 World of Open Source Survey, Q16 (n = 224, excl. 66 unsure), Q17 (n = 243, excl. 47 unsure), Europe only.

AI use patterns also risk undoing decades of hard-won data sovereignty progress, reintroducing exactly the dependency the first section of this chapter describes. And small projects see contributions drop as downstream users let coding assistants deploy and customise their software without ever contributing improvements back upstream.

Read together, the strongest agreement in the survey clusters around AI amplifying existing use and lowering the bar to contribute, while the weakest and most uncertain is maintenance automation. AI use produces two genuinely opposing effects at once: it amplifies existing open source use and lowers the bar to contribute, while at the same time overwhelming maintainers and developers, undermining community cohesion and disrupting the talent pipeline the profession depends on.

Which side will dominate is the question this precarious new balance leaves open. Across the interviews conducted for this report, opinion is cautiously optimistic more often than not, and the optimistic case rests on three expectations: that the positive effects will outweigh the disruption to community collaboration patterns, that every AI-detected security vulnerability fixed upstream raises that project's overall security posture, and that junior and senior developers alike will adapt as talent develops accordingly, even though the disruption is already visible today.

Interviewed experts mostly agree that the high-water mark of incoming security reports and AI-slop contributions has not been reached yet. Proactive upskilling and talent development will help organisations and individuals cope with what is still to come.⁴

The talent and pipeline question cuts in two directions at once. If AI removes the software-production bottleneck, one of open source's traditional strengths—a low-cost training ground for junior developers—is undercut along with it. Open source projects may in turn become the new proving ground precisely because the commercial pressure to reach for AI first is largely absent there. The same 2026 State of Tech Talent Europe report cited above shows a net 27 percent AI-driven hiring effect for Europe overall, positive in aggregate and not the squeeze the pipeline concern above worries about. The real strain is narrower and more specific: cybersecurity roles are understaffed at 48 percent of organisations, 14 points worse than the rest of the world, and organisations are 3.7 times more likely to upskill existing staff than to hire externally to close the gap. The net effect is that hiring is not contracting so much as reallocating, away from broad junior pipelines and toward upskilling people already inside the organisation. An open source track record already functions as a hiring signal in its own right, specifically for AI-specialist candidates.⁵

⁴Marco Gerosa and Adrienn Lawson, "2026 State of Tech Talent Europe: AI, technical hiring, and the skills gap in Europe," foreword by Thierry Carrez, The Linux Foundation, June 2026, <https://www.linuxfoundation.org/research/tech-talent-europe-2026>

⁵Marco Gerosa and Adrienn Lawson, "2026 State of Tech Talent Europe: AI, technical hiring, and the skills gap in Europe," foreword by Thierry Carrez, The Linux Foundation, June 2026, <https://www.linuxfoundation.org/research/tech-talent-europe-2026>

A licensing and provenance shadow hangs over all of this too, recurring across nearly every interview. Legal clarity on AI training data could still take years to settle; licensing and copyright remain, in one interviewee's phrase, "Licht und Schatten," light and shadow together, not a solved problem. And AI slop is now common enough that projects are building tooling just to measure how much of a given repository is AI-generated in the first place.

Project-level policies on AI-assisted code diverge sharply. No specific policy has yet emerged as an open source community best practice. One posture, associated with the Linux kernel, holds the human submitting a patch fully responsible for it regardless of tooling: use whatever helps, as long as the result meets the same review and provenance requirements, Developer Certificate of Origin included, as any other contribution. A more cautious posture, taken by some projects, restricts or declines AI-generated contributions outright rather than accept the copyright and provenance uncertainty that comes with them. Both are rational responses to the same underlying uncertainty, who is accountable for the code and what its copyright status actually is, reflecting different risk tolerances rather than one side being right and the other wrong.

AI and cybersecurity

As of this moment in September 2026, the OpenStack security team has issued 38 advisories for the year, roughly thirty times the pre-AI pace.⁶ The surge comes in two distinct tracks. One is a genuine rise in valid findings: AI-assisted scanning surfaces defects that would previously have gone unnoticed for years, and someone still has to triage, patch, and test each one before a fix can ship with a clear conscience. The other is noise, a parallel rise in low-effort, often AI-generated vulnerability reports that turn

out not to matter. Some communities simply dismiss these quickly, at real but comparatively modest cost; other projects, by several accounts, have been hit harder.

Whether the surge tapers is an open question. One working theory is that AI-assisted tools are working through a specific, finite class of defects, an exhaustible backlog, so the pace of discovery should slow once that class is largely cleared. However, if the underlying models keep improving, they may simply start finding new classes of weakness the current generation cannot see. In this case the discovery rate never really slows, it just resets against a harder problem. New code being created by humans and robots will also introduce new security vulnerabilities. Nine months into the year, the statistics available do not yet show deceleration either way.⁷

Somebody still has to maintain the overall open source components and fix the findings. At the same time, cybersecurity is the most acutely understaffed specialism in Europe's AI-era hiring reallocation. The overwhelming rise in advisories lands on whatever capacity remains. Responsible disclosure, the practice of giving maintainers a fixed window to patch before a vulnerability becomes public, was designed for a world producing findings at a much slower rate; several accounts describe that process now under real strain, not because the process is wrong, but because the volume it was built for no longer matches the volume it receives.

Regulation compounds the pressure at exactly the point where the stakes are highest. This chapter's first section already noted that the EU AI Act's open source exemption evaporates for high-risk systems: no carve-out at all, full conformity assessment and CE marking required from 2 August 2026. Contributors to open foundation models for the energy grid frame the underlying issue as one of calibration rather than compliance paperwork: a spelling mistake in an AI-drafted email is a minor annoyance. The

⁶ OpenStack, OpenStack Security Advisories, <https://security.openstack.org/ossalist.html>

⁷ Sonatype, The Hidden Dependency Gap in AI Agents, <https://www.sonatype.com/state-of-the-software-supply-chain/2026/ai-agents>

impacts of a single wrong AI-assisted decision in grid control will be noticeable in real life. The tolerance for error that makes AI broadly useful in low-stakes settings does not automatically transfer to critical infrastructure.

This chapter's first section also noted that AI models can leak data users supply during ordinary use back into their own training. The operational response several interviewees describe is independent of whether the model in question is open or proprietary: consider running models locally, and do not upload anything you cannot afford to see resurface somewhere else. Regulated sectors appear to take this more seriously. The financial industry, subject to audit requirements that put a compliance seat directly on the board, was singled out as doing this well, running models locally, accepting that the output is probabilistic and non-reproducible, and scoping its use accordingly. Healthcare, by the same account, is a more mixed picture: some practitioners understand the data-provenance risk clearly, others reach for a shortcut without asking where the data goes.


None of this is new in kind. Concerns about the security of open source components were already the single most cited factor limiting organisations' use of open source before any of this AI-driven volume arrived, cited by 39 percent of respondents, ahead of licensing concerns and lack of technical support (Figure 6). AI makes this concern more acute, with more signal, more noise, and more urgency. The regulatory response to that shift, the Cyber Resilience Act's compliance clock, its steward-status carve-out for open source projects, and the cost that clock quietly transfers onto maintainers, is a large enough story that this report addresses it separately, in the chapter on digital sovereignty and regulation.

Open source licensing and open source collaboration are not disappearing in the AI era. Their role is changing. There are good reasons to think open source digital infrastructure becomes even more important. Maintaining software and analysing its

FIGURE 6

Security concerns are the main limiting factor for OSS use

Which of the following factors limit OSS use in your organisation?



Source: 2026 World of Open Source Survey, Q24, n = 242, Europe only, excl. 14 unsure.

security in the presence of AI-generated code and AI-discovered vulnerabilities is expensive, and understanding, integrating and fixing what AI surfaces still requires real people. Open source infrastructure, Linux, the CNCF's portfolio of projects, Yocto, Zephyr among others, already receive significantly more sustained attention and pooled security analysis than the average proprietary product. This head start should compound as AI tooling matures. If anything, open source software is best placed among all software to benefit from what AI newly makes possible, precisely because its development already happens in the open, where AI tools can observe, learn from and contribute to it. The current upheaval remains genuinely painful for the people doing the work: overburdened maintainers were a problem well before AI arrived, and AI has made that problem worse. But the overall trajectory looks cautiously positive: more contributions, an improved security posture, lower barriers to entry, provided the ecosystem finds a way to carry its contributors and maintainers through this transition.

Community and technical governance


Half of European organisations that fork open source components privately cannot say how many hours a release cycle costs them. Nearly a third report no formal open source governance at all (Figure 7). And the single most cited blocker to adopting open source in the first place, security, is cited by only a few points more of the sample than the fourth-ranked blocker, a simple

lack of leadership buy-in. Individually, each of these findings describes a separate gap. Together, they describe an ecosystem in which organisations use open source extensively, contribute to it unevenly, and manage the resulting risk and opportunity only where someone inside the organisation happens to have both the mandate and the visibility to do so.⁸

FIGURE 7

Europe trails behind the rest of the world in open source governance maturity

Which of the following best describes your organisation's approach to open source governance?



Source: 2026 World of Open Source Survey, Q12. Europe n = 339 (excludes 28 don't know); Rest of World n = 469 (excludes 25). Stages follow the TODO Group's five-stage OSPO maturity model (github.com/todogroup/ospology). Markers show Europe minus the rest of the World in percentage points, calculated before rounding.

This chapter follows that thread through four questions, narrowing from the structural to the individual. How mature is European governance of open source, and how does maturity benefit organisations? Does contributing upstream produce real influence over the projects an organisation depends on, or is influence earned some other way? What does the alternative to contributing, forking and maintaining privately, actually cost, now that two independent studies have put a number on it? And when adoption fails to happen at all, what stops it, a shortfall in the technology, or a shortfall in the process built to evaluate it?

The short answer, developed across the four sections that follow, is that governance, contribution, forking and adoption are not four separate problems but four different views of the same organisational and procedural capacity or capacity gap.

⁸Adrienn Lawson and Sam Boysel, "ROI for Open Source Software Contribution: Insight from the Open Source ROI Survey and Economic Model," forewords by Chris Aniszczuk, Hillarie Prestopine, and Masato Endo, The Linux Foundation, February 2026, <https://www.linuxfoundation.org/research/contribution-roi>

Governance maturity

Not one of the 31 European Union-headquartered public-sector organisations in this survey's sample reports "compliance and inventory," the first stage of the TODO Group's five-stage OSPO maturity model, as its governance stage. The plurality instead reports no formal governance at all: 15 of the 31 respondents, or 48 percent, roughly half the European sample overall. Across Europe as a whole, governance sorts into no formal governance (31%), compliance and inventory (32%), community education (12%), active project leadership (17%) and strategic leadership (8%) (Figure 7), a distribution that already lags the rest of the world at both ends, 12 points more likely to have no formal governance at all and 6 points less likely to reach strategic leadership.

Interviewees offer multiple overlapping explanations for why maturity clusters this way. A mindset gap treats open source as a free input to be consumed rather than a strategic lever to be

managed. A capital-structure mismatch leaves Europe's more fragmented market unable to accumulate the governance weight that a single large domestic platform firm could. A capability gap means no staff are assigned to community roles at all, or only in an operational capacity. Procurement path dependency and decades of outsourcing IT has eroded the in-house understanding organisations would need to engage upstream even if they wanted to. A competing principal, a ministry, a dominant funder, or a US parent company sets direction and leaves no mandate for anyone to represent the organisation's interests in a project's governance. And a value-chain blindness, a cultural default to treating software as a purchasable product rather than a supply chain one participates in, is reinforced by procurement structures that were themselves built around that same default.


Governance maturity is not a side variable. It correlates with almost everything else this report examines. Organisations with no formal governance make upstream contributions 30 percent

of the time; organisations with advanced governance do so 87 percent of the time. The same climb holds for forking (29 percent to 75 percent) and for participating in project governance directly (18 percent to 78 percent) (Figure 8). The one engagement mode that runs the other way is passive, unmodified use of components as-is, which is highest at basic maturity (91 percent) and lowest among advanced organisations (82 percent): the more governance capacity an organisation builds, the less it simply consumes open source without touching it, and the more it forks, contributes and shows up to decide.

FIGURE 8

Governance maturity correlates with deeper engagement with open source

In what ways does your organisation typically engage with open source software? – by governance maturity



Source: 2026 World of Open Source Survey, Q13 × Q12, Europe (n = 339). Each dot shows the share of that governance level answering yes, with don't-know responses excluded from its base.

That deeper engagement pays for itself. Organisations with no formal governance report a mean benefit-to-cost multiple of 3.6 times from their open source use; organisations with advanced governance report 4.6 times, against an overall European mean of 4.1 times (Figure 9). Dependency visibility follows the same gradient, if less steeply: only 15 percent of organisations with no formal governance say they fully know which open source components sit in their products or services, rising to 25 percent

at advanced maturity (Figure 10). None of this proves governance causes the higher return rather than simply correlating with organisations that were already more capable in other ways. But the direction is consistent across every measure this survey took, and it argues against treating governance investment as a compliance cost that competes with the benefits open source delivers, rather than as one of the mechanisms that produces them.

FIGURE 9

Organisations with stronger governance report higher returns from open source




Source: 2026 World of Open Source Survey, Q23 x Q12, averages, Europe only, n = 87.

Smaller organisations illustrate the tension most sharply. SMEs report no formal governance at nearly double the rate of larger organisations, 44 percent against 23 percent, and make upstream contributions less often, 48 percent against 64 percent (Figure 11). Yet the same SMEs capture value from open source far out of proportion to their size: annual savings equivalent to 37.5 percent of revenue, against 1.2 percent for larger organisations (Figure 12), and they report every major benefit, lower vendor lock-in,

FIGURE 10

Better governance correlates with better dependency visibility

Do you know which open source components and dependencies are present in your products or services?




Source: 2026 World of Open Source Survey, Q28 x Q12, Europe. Overall donut: n = 231 (excludes 10 unsure). Segmentation: n = 221; each bar shows the share answering fully within that governance level, excluding unsure.

improved productivity, and lower cost of ownership, more often than larger peers do. Small organisations with the least capacity to formalise governance are, by this evidence, the ones with the most to gain from doing so. Governance maturity here looks less like a luxury only large firms can afford than like an underused opportunity smaller organisations have not yet been shown how to pull.

FIGURE 11


SMEs lag in maturity




Source: 2026 World of Open Source Survey, Europe. Q12 n = 316 (106 SMEs, 210 larger); Q13 n = 296 (104, 192), excluding unsure. SMEs have fewer than 250 employees and under \$50 million in revenue (Q10, Q11).

FIGURE 12

SMEs capture greater value from open source relative to their size



AND THEY REPORT THE BENEFITS MORE OFTEN



Source: 2026 World of Open Source Survey, Europe. Savings as a share of revenue: Q22 ÷ Q11 at response-band midpoints, median organisation; n = 52 SMEs and 84 larger. Benefits: Q21, share answering often or always; n = 71-73 and 150-154, excluding N/A by item.


Where organisations do engage with governance, they show a clear institutional preference and a due-diligence gap. Foundation-hosted, neutrally-governed projects are the model European organisations are most likely to adopt (at 52 percent), 8 points ahead of the rest of the world (Figure 13), ahead of both company-backed and purely community-maintained alternatives. But the process for getting there is thinner than elsewhere: only 27 percent of European organisations apply a formal, consistently

applied process to evaluate a project’s health and governance before adopting it, 13 points below the rest of the world, while 36 percent rely on informal checks by individual developers, 8 points above (Figure 14). Europe prefers openly and well-governed projects in principle while relying, in practice, on others to govern them.

FIGURE 13

Foundation-hosted projects are the preferred governance model

Which open source governance model makes your organisation most likely to adopt a project?




Source: 2026 World of Open Source Survey, Q31. Europe n = 143 (excludes 43 unsure and 55 with no preference); Rest of World n = 233 (excludes 45 and 62).

FIGURE 14

Europe relies more on informal project evaluation than the rest of the world

Before adopting an open source dependency, how does your organisation evaluate the health and governance of the project?



Source: 2026 World of Open Source Survey, Q30. Europe n = 217 (excludes 24 don't know); Rest of World n = 306 (excludes 34 don't know).

The throughline is not that governance is a virtue to be preached at organisations that lack the resources for it. It is that governance functions, empirically, as a multiplier: on engagement depth, on financial return, on dependency visibility, and, as later sections of this report set out, on regulatory readiness as well.


An organisation's governance stage is less a description of how mature it has become than a fairly reliable predictor of how much more it is likely to get out of everything else it already does with open source.

Does contributing upstream actually buy influence?

Contributors who ask do better than those who do not. Among European organisations that contribute upstream, 83 percent get warning before a breaking change lands. 69 percent say contributing has made their own development faster (Figure 15). 62 percent who try to influence a project’s roadmap succeed at least half the time. 53 percent get quicker answers from maintainers. Being social processes, none of these approaches comes with a guarantee, but together they describe a mechanism that works often enough to be worth investing in and using deliberately.


Interviewees are consistent on how that mechanism works: contribution earns influence, money alone does not. Sponsoring a conference or paying a foundation membership fee does not move a project’s direction; showing up on every call, writing the pull requests, and doing the unglamorous maintenance work does. One interviewee draws the comparison to standards bodies directly: whoever is present for every meeting and does the drafting sets the agenda, whether the venue is a standards committee or an open source project’s governance list. When asked whether contributing upstream has bought their own organisation real influence, contributors in the local survey answer roughly as this mechanism would predict: about half say sometimes or frequently, against three in ten who say rarely, which is itself consistent with influence being earned case by case rather than conferred automatically by the act of contributing at all (Figure 16).

FIGURE 15
Contributing upstream creates value beyond code



Source: 2026 World of Open Source Survey, Europe. Q40 n = 105; Q37 n = 99; Q38 n = 100; Q39 n = 92. Asked only of organisations contributing upstream (Q13), excluding unsure. Q38 is rebased on those who try to influence the roadmap.

FIGURE 16
Upstream contributions’ influence on project roadmaps



2026 World of Open Source Survey, Q38, n = 111, Europe only, excl. 14 unsure, asked only of upstream contributors (Q13).

That the mechanism works does not mean Europe uses it. One interviewee puts European organisations' technical contribution to a major cloud-native foundation at a quarter to a third of the total contributions to the foundation, a substantial share by any measure, while European presence at that foundation's governance level remains close to absent. The gap is not a lack of two routes to influence, election on technical merit or a fee-based board seat, so much as a failure to take either route systematically. As the previous section's governance figures already show, only 37 percent of European organisations participate in project governance or decision-making at all (Figure 17); contributing code without ever showing up to the meetings where direction gets decided leaves exactly the kind of influence this section describes on the table.

Not every organisation reaches influence through contribution at all. One public-sector organisation requires, by contract, that


every feature it commissions from a vendor be published to the community edition of the underlying project, an indirect but real route to influence that never shows up as a contribution in the usual sense. Another turned an internal calculation engine into the open source standard used across an entire European infrastructure sector simply by open-sourcing it early and well: other organisations adopted it, universities built research on it, and the standard-setting itself became the influence, achieved without a single upstream pull request to somebody else's project.

The alternative to contributing has its own cost, one this report examines next: an organisation that forks rather than contributes still has to maintain what it forked, indefinitely, on its own, and that cost is considerably easier to measure than the influence a contribution might have bought instead.

FIGURE 17

Many organisations adopt open source and contribute upstream, but relatively few participate in project governance

In what ways does your organisation typically engage with open source software?



Source: 2026 World of Open Source Survey, Q13. Europe n = 312–350 (excludes 17–55 unsure); Rest of World n = 406–461 (excludes 33–88 unsure).

Framed as an investment, contribution pays for itself at a rate few procurement processes would recognise as a return: European organisations that make code contributions report a mean benefit-to-cost multiple of 3.3 times, community contributions 2.5 times, and direct financial contributions to projects or foundations 2.3 times (Figure 18). Set against the earlier finding that lack of leadership buy-in is the factor European public-sector organisations cite most often as a blocker to open source adoption, the pattern looks less like contribution failing to pay off and more like too few organisations being positioned, or permitted, to try it at all.

The real cost of maintaining private forks

Nearly half of European organisations, 48 percent, maintain private forks of open source components rather than consuming and contributing back to the upstream original. The typical forking organisation spends 311 hours per release cycle keeping its forks patched, nine and a half forks on average, at roughly 33 hours of maintenance labour each. The mean organisation, pulled upward by a long tail of heavy forkers averaging 91 forks apiece, spends nearly ten times that, close to 3,000 hours per cycle (Figure 19). Both figures describe the same underlying activity; they diverge because a small number of organisations fork far more heavily than the typical one.

Even organisations that maintain forks struggle to say how much doing so actually costs them. Asked directly how many hours their forks require per release, the largest single response among contributor-gated

FIGURE 18

Contributions pay for themselves



Source: 2026 World of Open Source Survey, Q41, Europe, averages. n = 57 for code, 56 for community and 54 for financial contributions, excluding unsure. Asked only of organisations contributing upstream (Q13). Respondents estimated the multiple by which contribution benefits exceed costs; for example, 2 means benefits are twice costs.

FIGURE 19

Private forks create maintenance burden



Source: 2026 World of Open Source Survey, Q13 (n = 330), Q32 (n = 102), Q33 (n = 73), Europe. The products combine three separate measurements; they are not costs measured at any one organisation. Hours estimated from response-band midpoints, with the open top band at 1.5× its lower bound. Unsure responses excluded from each base.

respondents to the local survey is “don’t know or not sure,” 28 percent, ahead of every specific hours band on offer. This is somewhat of a counterintuitive finding. Organisations that already contribute upstream, and so have some occasion to think about the alternative, still cannot put a number on what forking costs them.

Interviewees describe three reasons the number stays untracked. Forking is recognised, in principle, as more expensive than upstreaming, several say so explicitly, but the comparison is rarely made with real figures on either side; it is a belief held with confidence and measured with none. Public-sector organisations in particular describe a capability gap between deploying open source and actually maintaining the customisations layered on top of it: one interviewee’s organisation began with external contractors before hiring its own DevOps staff internally, having realised it had, in its own words, too little control over what it was running. And fragmented, short-term public funding compounds the problem structurally: EU and national programmes that fund three to five year projects, then stop, leave behind patches and customisations nobody has a mandate or a budget to upstream once the funding ends. One interviewee has a word for the result, published to a code repository and then abandoned, that translates roughly as “zombie files”: present, discoverable, and never touched again.

This has a regulatory dimension the report addresses fully elsewhere: staying on an upstream project’s security-patch feed is one core compliance ask behind the Cyber Resilience Act. Every hour spent maintaining a private fork instead is, in effect, a direct cost on meeting it.

Set against what contributing upstream returns—3.3 times benefits to costs for code contributions, 2.5 times for community contributions (Figure 18)—on the evidence set out in the previous section, forking looks like the more expensive choice by a wide margin. Yet it remains the default for nearly half of European

organisations (Figure 19). The gap between what the arithmetic favours and what organisations actually do is not, on this evidence, a gap in incentives. It is a gap in visibility: an organisation that has never measured what its forks cost has no number to compare against what contributing would have cost instead, and so keeps making the more expensive choice.

What actually blocks open source adoption


Ask European organisations what limits their use of open source and no single factor dominates. As observed in Figure 6 above, security concerns lead, but only just, at 39 percent, ahead of licensing and IP concerns at 36 percent, lack of technical support at 35 percent, lack of leadership buy-in at 34 percent, and lack of understanding at 34 percent. That flat distribution masks a sharp split by sector. Lack of leadership buy-in is cited by 55 percent of EU-headquartered public-sector respondents, against 34 percent overall, the single largest segment gap in this part of the survey. It is also the figure that best explains why interviewees describe this problem in structural and procedural terms rather than technical ones. Procurement structure comes up repeatedly: tenders built around bundled proprietary licences and services simply do not fit an open source licence paired with a separate support contract, so open source offers lose by default rather than on merit. Missing expertise compounds it, procurement officials are rarely trained to evaluate what an open source offer actually is, and IT decision-makers do not always understand its strategic value well enough to make the case to the leadership whose buy-in the survey shows is missing. One interviewee’s proposed fix is blunt: separate technology procurement from service procurement, and certify procurement officials the way other professional judgements already require certification.

A related but distinct question asks organisations that already contribute what specifically limits their contributions, rather than

their adoption in the first place, and the answer there looks different: legal or licensing concerns, 42 percent, ahead of a fear of leaking intellectual property, 40 percent, and uncertain return on investment, 33 percent (Figure 20). The IP-leak concern rose nine points against the equivalent 2025 figure, the sharpest year-over-year move in that block, while uncertain ROI eased slightly. Adoption and contribution, in other words, are blocked by overlapping but not identical fears: one is about whether to bring open source in at all, the other about whether to give anything back once it is already there.

FIGURE 20
Governance concerns constrain contributions

Which factors limit OSS contributions in your organisation?



Two further beliefs recur across interviews without a matching survey item to test them against. A myth of insecurity, open source characterised as built in a cellar, persists in some regulated sectors despite evidence to the contrary. The open source and cybersecurity communities are not always well aligned in correcting it. An IP secrecy belief, the conviction among compliance and legal teams that publishing source code incurs unknown risks and exposes trade secrets, generates review and approval cycles lasting months at organisations trying to open even a single repository. These cycles exist to disprove a belief

rather than to manage a real risk. A third, an awareness gap, is narrower still: some public-sector buyers simply do not know that open alternatives to the incumbent vendors exist, one interviewee compares it to only ever buying cars from a single country because nobody has explained why any other option is worth considering, and default to the familiar not out of preference but out of not knowing there was a choice to make.

Procurement is also where interviewees see the most concrete lever for change. European public-sector spending is enormous in aggregate, but fragmented across 27 national markets in a way that dilutes its influence as a single buyer. Federal and national software licensing budgets, running into hundreds of millions of euros annually in at least one country cited, are large enough to redirect entire markets if procurement criteria required it. One emerging criterion several interviewees point to is demonstrable portability between infrastructure providers, sometimes called switchability, already met in practice by at least one major public-sector software platform certified to migrate between competing cloud providers. Where governance-model preference was already uneven earlier in this chapter, foundation-hosted projects the clear plurality choice overall, EU-Public respondents split their preference far more evenly between foundation-hosted, company-backed and no preference at all. As a result, procurement criteria do not clearly favour open technologies.

The technology works. Every blocker catalogued in this section is a mismatch in how organisations evaluate, procure or govern open source, not a shortfall in what open source itself delivers: procurement processes not built to recognise it, staff not trained to evaluate it, leadership not yet backing it, and beliefs about its risk that survive independently of the evidence. That is, on balance, the better diagnosis to have: organisational and procedural barriers can be redesigned by the organisations that hold them, in a way that a genuine technical shortfall could not be. Procurement, in particular, has real room to catch up with a technology whose case is already made.

Digital sovereignty and technological resilience

Digital sovereignty has firmly dominated the global technology debate over the past year. In her 2025 State of the Union address, European Commission President Ursula von der Leyen called for Europe to assert its technological independence, framing digital autonomy as a prerequisite for long-term economic resilience and governance⁹. In her 2026 speech setting the course for the European Union, she put emphasis on the impact of AI, marking the importance of cooperation and guardrails under the AI Act.¹⁰ The European approach, centered on human agency, responsibility, and energy efficiency, was strongly highlighted. She also announced upcoming initiatives to accelerate industrial AI deployment across five high-value sectors: health, transport, agri-food, advanced manufacturing, and defense and space.

In Europe, digital sovereignty has transitioned from a policy concept into an immediate operational priority, reflected in new policy files introduced in 2026. Findings from our 2026 survey reveal that similarly to the rest of the world, 94% of European organisations consider digital sovereignty important to their technology strategy, with sixty-one percent rating it as very important (Figure 21). The definitions and understanding of digital sovereignty vary across institutions and organisations, and especially across global geographies, but it is clear that open source has resurfaced as one of the key topics linked to reducing critical dependencies in the digital sphere. EU digital sovereignty policy is trying to find a delicate balance between reducing foreign dependencies on one side, and avoiding isolation and tech fragmentation on the other. Open source has been recognised

as a way to strike that balance, which is why it is emerging at the center of policy initiatives this year.


FIGURE 21

Digital sovereignty has become a strategic priority

How important is digital sovereignty to your organisation's technology strategy?

94% of European organisations call digital sovereignty important to their technology strategy.
61% call it very important.

What is driving digital sovereignty's importance in your organisation?



Source: 2026 World of Open Source Survey. Q18: Europe n = 264 (excludes 26 unsure); Rest of World n = 351 (excludes 36). Important combines very and somewhat important. Q19: Europe n = 229 (excludes 3 unsure); Rest of World n = 326 (excludes 4). Select all that apply; shares do not sum to 100%. Differences calculated before rounding.

⁹State of the Union address by Ursula von der Leyen, 2025 https://commission.europa.eu/strategy-and-policy/state-union-addresses/state-union-2025_en

¹⁰State of the Union address by Ursula von der Leyen, 2026 https://commission.europa.eu/strategy-and-policy/state-union_en

We can easily identify two main digital sovereignty themes weaving through responses to our survey and expert interviews, as well as appearing in several new EU policy proposals relevant to open source unveiled in 2026. The first one is the increasing focus on geographical location in the digital sphere. It is motivated by geopolitical tensions, diverging regulatory frameworks and political approaches between regions, increased competition, and rising anti-globalist sentiments around the world. The second theme is Europe's bet on public procurement as a demand-side mechanism for advancing bloc's competitiveness and decreasing its dependence on third-country providers and political regimes. While open source has been historically difficult to be addressed in public tenders, policymakers took big steps this year in order to fill this gap and to leverage open technology for increasing digital sovereignty.

“Tech sovereignty is Europe's ability to act independently in the digital world by developing and controlling key technologies, data, and infrastructure, while reducing reliance on non-EU providers.”¹¹

- European Commission

“Digital sovereignty describes the abilities and opportunities of individuals and institutions to be able to perform their role(s) in the digital world independently, self-determinedly and securely.”¹²

- Center for Digital Sovereignty (ZenDiS), Germany

Open source as a potential means for achieving digital sovereignty is present not only in discourse and policy, but mirrored across the global economy. The United States continues to reinforce its capacity through market-driven commercial open source investments estimated at \$7.7 billion annually.¹³ China has long leveraged open source as a cornerstone of state planning, embedding open technologies within its Five-Year Plans, early Golden Projects, and e-Government systems while navigating international security scrutiny and export restrictions. Emerging economies are taking similar strides: India has established its digital

backbone through the “India Stack” while nations across South America, such as Brazil, and Africa are actively building public digital infrastructure around open principles to safeguard strategic autonomy. On the global level, the UN is continuing and strengthening its involvement with open source and digital commons.

What are the drivers behind the push for digital sovereignty in Europe? Our survey shows that European organisations, similarly to global peers, prioritise security and privacy, cited by 71% of respondents (see Figure 21). While privacy considerations have been on the European agenda for years, cybersecurity concerns have grown at a faster pace more recently, motivated by geopolitical tensions, accelerated speed of discovery and emergence of vulnerabilities, as well as regulatory advancements and new obligations in the EU.

Mitigating vendor lock-in and cost risks comes second, noted by 55% of respondents in the EU compared to 48% for the rest of the world. The biggest differentiator in digital sovereignty drivers, twelve-percentage points higher in Europe compared to the rest of the world, is the advancement of national or regional interests, cited by 47% of respondents. This data might point to the growing realisation that European organisations are mostly reliant on third-country technologies, as well as the trend in the discourse towards developing native solutions and supporting local companies. While global organisations are more likely to view digital sovereignty as a mechanism for strategic or competitive advantage (50% rest of the world vs. 45% Europe) or stakeholder and investor interests (50% rest of the world vs. 43% Europe), European organisations treat open infrastructure primarily as a defensive shield against supply chain vulnerability, price volatility, and regulatory exposure.

¹¹European Commission, EU Tech Sovereignty Package proposal, 3 June 2026 <https://digital-strategy.ec.europa.eu/en/policies/eu-tech-sovereignty>

¹²Center for Digital Sovereignty (ZenDiS), Germany, <https://www.zendis.de/en/our-mission>


¹³Sam Boysel et al., 2024 Open Source Software Funding Report, <https://opensourcefundingreport2024.com/>

Surveyed organisations almost unanimously agree that open source software is effective in meeting technological independence demands (Figure 22). Success is reported in using open source for cloud and infrastructure control, application portability, data localisation, and interoperability. The perception of open source's usefulness for local AI deployment and custom AI systems is less optimistic, with only 33% of respondents saying it's very successful, covered in more depth in the Open source and AI sovereignty chapter of this report.

FIGURE 22

Open source is central to achieving digital sovereignty

In which areas has open source been successful in helping your organisation achieve its digital sovereignty goals?



Source: 2026 World of Open Source Survey, Q20, n = 133–188, Europe only, excl. 44–99 not applicable or not sure by row, asked only of those who rated digital sovereignty important in Q18.

In our 2025 report, we described an ecosystem in transition - open source adoption was widespread and has been consistently growing in recent years, but the recognition of open technologies

as a strategic imperative in Europe keeps increasing, as visible in our survey and expert interviews. Accompanying this progression, there have been voices concerned with isolationist framing that could fragment global collaboration, warning against pursuing notions of “European open source.” Software developed in regional isolation risks depriving Europe of global innovation, diluting public investments, and creating isolated silos. These concerns seem to have been recognised in European policy to some extent, as highlighted by the European Commission’s Executive Vice-President for Technological Sovereignty, Security, and Democracy, Henna Virkkunen, proclaiming that “Technological sovereignty does not mean protectionism.”¹⁴

The latest European governmental policies point to openly governed technologies as a cornerstone of their digital sovereignty strategy. The 2026 European Commission’s Tech Sovereignty Package, a landmark legislative proposal that might mark this year as the beginning of a new era of a set of EU laws, boldly bets on open source for digital sovereignty and is overall framed through the need for a vibrant open digital ecosystem in Europe.

Because the technologies underpinning cloud, AI, and modern infrastructure already exist as openly governed global commons, Europe’s fastest and most realistic path to technological sovereignty lies in actively participating in and influencing these ecosystems. Strategic upstream investment aligns with core European values of transparency, accountability, and interoperability while providing direct control over critical technology roadmaps. These values are clearly delineated in new policy proposals in the EU in 2026, marking a more decisive direction for both technological independence, and open source software in the continent.

¹⁴Remarks by Executive Vice-President Virkkunen and Commissioner Jørgensen on the tech sovereignty package, 3 June 2026 https://ec.europa.eu/commission/presscorner/detail/da/speech_26_1244

EU regulatory frameworks: the Tech Sovereignty Package

The European Commission published its much anticipated Tech Sovereignty Package in June 2026, a landmark policy proposal planned as key for increasing Europe's competitiveness, resilience, and strategic autonomy. It comprises four major policy actions, with two of them being the most pertinent to the open source ecosystem which are the Cloud and AI Development Act and the Open Source Strategy. It is the very first time that a major legislation proposal in the world bets so openly on open source software, with open development being often previously treated as a somewhat niche, deeply technical topic. This perception has changed, motivated by strong political and economic drivers. Europe relies on non-EU tech for over 80% of digital products, services, infrastructure, and IP, and the European public sector currently spends an estimated €264 billion annually on IT, overwhelmingly directed toward third-country proprietary vendors. Additionally, the digital landscape is changing at an accelerating pace, often too dynamic for lengthy and complicated legislative processes. 2026 can be marked as EU legislators' step towards getting up to speed with those developments, and an attempt to align with market reality, as well as recognising the growing maturity of the open source system.

Central to the Tech Sovereignty Package is the proposal for the Cloud and AI Development Act (CADA), which aims to unlock European cloud and AI capabilities and harmonise the approach for EU's public sector cloud procurement. It calls for a coordinated 'ecosystem approach', placing a special focus on open source as a lever for digital sovereignty and a way to reduce strategic reliance on third-country providers, and plans to at least triple the EU's data centre capacity within the next 5–7 years.

CADA proposes some crucial provisions on open source, including Article 41 setting the obligations for EU institutions and

public sector bodies to encourage, facilitate and prioritise open standards and open source components when building their cloud and AI stacks. Moreover, it legitimises the existing EU Open Source Software Catalogue and makes it mandatory for public sector bodies, alongside a coordinated network of Member States Open Source Programme Offices (OSPOs) which used to operate in a more informal capacity until now.

A key pillar of CADA is its harmonised cloud computing sovereignty framework for public services, establishing a 4-level assurance scheme. Cloud computing service providers would need to perform a conformity self-assessment (level 1) or undergo an independent audit for levels 2–4, with level 4 being reserved for the most sensitive critical infrastructure with full supply chain sovereignty. It will be EU Member States designating which levels their public bodies' activities require, in line with NIS2. This progressive assurance ladder, in its current shape, would designate levels based on criteria such as EU ownership of a provider, if technical and security personnel are EU nationals/ residents, data being processed under EU jurisdiction, and transparency over the software and hardware bill of materials. Moreover, levels 2–4 require cloud service providers to submit a comprehensive SBOM mapping of all dependencies, with details on reliance on non-EU software or open source software. This regulation will undergo rounds of negotiations and is planned to enter into force the earliest in 2027/2028.

The second policy initiative that is part of the Tech Sovereignty Package and will influence the open source ecosystem in Europe is the EU Open Source Strategy, which positions open software as a foundational pillar across public institutions and enterprise stacks in order to foster sovereignty, competitiveness and security. It is a strategic Communication, taking immediate effect to guide public sector procurement, internal Commission practices, and Member State alignment, and has been overall well-received by open source communities, especially that it expands the scope of the previous, internally focused 2020–2023 strategy to the broader

European digital economy, in line with the ecosystem approach of the whole Package.

In order to gather stakeholder feedback for the Strategy, the European Commission launched its Open Digital Ecosystems consultation in early 2026.¹⁵ The initiative became one of the Commission's most successful consultations, receiving over 1,600 responses from developers, public administrations, industry leaders, and research institutions. This unprecedented turnout demonstrates that open source is no longer viewed as a niche technical choice, but as a central economic and strategic priority across Europe's public and private sectors.

Planned measures include scaling the European Commission's Open Internet Stack (formerly Next Generation Internet initiative) into a single hub for open source building blocks; setting up European Digital Infrastructure Stewardship Organisation and stimulating the emergence of new European associations and foundations; providing tools to assess sovereignty of value chains in the EU; promote a variety of security measures for OS; partnering with EU Member States through Digital Commons EDIC; strengthening open source social media; establishing business accelerators for open source startups; supporting open source stewardship; and promoting open source in R&D including in Horizon Europe research programs. The Commission also plans to launch an Open Source Maintenance Instrument to support critical components or create European capacity to fork projects when necessary. To track long-term execution, it will establish a monitoring framework that requires annual reviews with Member States through the Digital Decade Board and a formal report submitted to the European Parliament every three years. In terms of investments, Europe is planning to mobilise an estimated €2 billion over the next seven years through the public and private sectors to operationalise the Strategy.

These are just some of the elements of the Strategy that is currently being internally prepared for implementation throughout the European institutions and EU Member States. Its goals are ambitious, but it is yet to be seen how it will be implemented throughout the Union.

"Technological sovereignty is thus grounded in openness, partnership and fair competition.

It does not mean isolation, protectionism, or tech decoupling. On the contrary, a technologically sovereign EU will continue to contribute to and benefit from integration in the global economy but will be better equipped to manage both the risks and the opportunities stemming from global technological interdependence."¹⁶ Open Source Strategy

While the Strategy openly rejects protectionism and isolationism, some of the provisions of the Tech Sovereignty Package that focus on geographical location raise questions in the open source community. The Package builds its stewardship and infrastructure around new EU-specific vehicles rather than existing global foundations and communities, which can provide time-tested governance practices and partnerships. The Communication announcing the Strategy states that "Open source building blocks are mostly maintained through foundations with US (and Chinese) big tech providing a massive share of the funding and contributions." While technically correct, it fails to note that in open governance, financial sponsorship is decoupled from technical leadership, and both are highly distributed. Financial sponsorship of a project does not give control over the technical artifacts to either the financial sponsors or the foundation itself, and geographical location of a steward does not predetermine where the contributors to a project are located.

¹⁵European Commission, Open Digital Ecosystems consultation, 2026 https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/16213-European-Open-Digital-Ecosystems_en

¹⁶European Commission, The EU Open Source Strategy, 3 June 2026, <https://digital-strategy.ec.europa.eu/en/policies/open-source-strategy>

European developers and maintainers drive a significant share of innovation in global, critical open source infrastructure software like OpenStack or Kubernetes with almost 40% of global contributions coming from Europe. Further strengthening European capacity and participation in the global commons would be more suitable to fulfill the digital sovereignty goals. Linux Foundation Research shows that active open source contribution delivers a 2 to 5 times return on investment, whereas passive consumption increases costly technical debt.¹⁷ Moreover, open source licensing, and the right to fork it preserves, means an EU-based open source steward is not a necessity in order to guarantee continuity of access to open source technology. What is needed is familiarity with open governance processes to reduce the hypothetical transition timeframe and to innovate and build on top of existing artifacts.

Public sector and procurement

The EU's public sector has been increasingly focusing on increasing its participation in open source throughout the past couple years, especially among the founding members of the Digital Commons European Digital Infrastructure Consortium (DC-EDIC): France, Germany, Italy, Luxembourg and the Netherlands. This initiative plans to pool resources and political commitment in order to develop and scale sovereign digital infrastructure in Europe. Unsurprisingly, public procurement is on top of the agenda for the group, just like for many EU Member States.

In the past year we saw some countries engage with open source meaningfully for the first time, such as one of EDIC's observers, Poland, that set up a new ministerial team and an OSPO, and places open technologies at heart of its digitalisation strategy. At

the same time, other EU Member States continue their mission, such as Germany through its Center for Digital Sovereignty in Germany and France through DINUM for a variety of projects around developing an open source stack responding to Europe's public sector needs.

Public procurement accounts for around fifteen percent of EU GDP - around €2.5 trillion per year - making it one of the most effective demand-side mechanisms to advance European digital sovereignty. In order to tap into this potential, the European Union started the process of revising its Public Procurement Directives in 2026, following public consultations in late 2025 - early 2026. The proposal for the new Public Procurement Act was published on 9 September 2026, aiming at harmonising procurement frameworks, and in the digital domain, at aligning with the planned provisions of the Tech Sovereignty Package.

Until now, the EU's procurement system was fragmented and favored price-only selection criteria. The new proposal moves towards best price-quality as a default, recognising the market-creating power of public procurement and favoring security and innovation. Open source is introduced in the context of innovation, explicitly allowing public buyers to consider purchase of open source solutions with an objective of "digital transformation and technological advancement including interoperability of and open source solutions for digital systems and elements of public services". While it could be seen as a win by the open source ecosystem, placing open technology primarily in the context of innovation is not consistent with the current market reality and could undermine its potential as a real alternative to proprietary solutions in procurement.

The proposal also allows public buyers to apply European preference requirements by limiting tender eligibility, requiring minimum Union or covered origin, or granting evaluation

¹⁷Adrienn Lawson and Sam Boysel, "ROI for Open Source Software Contribution: Insight from the Open Source ROI Survey and Economic Model," forewords by Chris Aniszczuk, Hillarie Prestopine, and Masato Endo, The Linux Foundation, February 2026.

preferences. It explicitly states that it's been designed to be a horizontal framework for European preference requirements to be referenced in the future. Previously, "buy European" rules were narrow exceptions in specific sectors and cases, and the new proposal provides a general tool available across standard public purchasing whenever strategic EU interests or market reciprocity are at stake.

What could it mean for open source? Many open source stakeholders and companies called for a more explicit inclusion of open source in the proposed Public Procurement Act than the one presented by the Commission. The proposal misses a chance to establish a rebuttable presumption favoring open source solutions whenever functional equivalence exists and to introduce a public money, public code principle. Linux Foundation Europe recommended these provisions in its response to the public consultation, alongside assessing software risk with objective open standards instead of origin-based criteria, focusing tenders around established open source standards rather than competing software builds, and favouring vendors who contribute upstream in order to stimulate a healthy open source ecosystem.

On the other hand, the Tech Sovereignty Package comprises several provisions recognising the role of open source in digital markets and will provide guidelines for public buyers to overcome structural and administrative hurdles in procuring open solutions. There is a high potential that this practical implementation support could have a larger impact on leveraging the power of procurement for reducing vendor lock-in and increasing technological sovereignty through open source for the EU than a large, general (not focused on technology) law that mandates buying open. At the same time, it is crucial that the European preference does not hinder procurement of solutions based on globally developed and maintained open source projects. There is still a long way to go until the proposal takes final shape, and it will include many forces at play, as it is a horizontal regulation touching a very large part of the EU economy and where stakes

are high for all sides involved.

Independently of EU policy developments, public administrations across the European Union increasingly seek to procure open source software in line with the strategic objectives of digital resilience. Public procurement rules should enable them to fulfill their goals, thus open source solutions should be referenceable in EU legislation and public procurement procedures under clearly defined requirements such as open governance and transparent development processes.

Other policy developments in 2026

Through vibrant and long-lasting consultations, policy engagement, advocacy and research, the CRA ended up introducing new roles of stewards and manufacturers for organisations handling open source components, while not introducing burden for individual contributors. The law revolutionised the approach to managing software supply chain risks and responsibility for digital products - and the effects are still only to be fully seen, as of the time of writing of this report, 11 September 2026, the reporting obligations of the CRA are becoming mandatory. The regulation will fully apply only at the end of 2027.

The CRA has shown how European regulators can work with the open source ecosystem in order to find a way of creating laws that fulfill institutional goals, while at the same manage to recognise the unique nature of open source. One of our experts points out that despite flawed post-negotiation adjustments, the regulation successfully leverages European market power to force software vendors into upgrading their security posture globally, while another perceives CRA obligations as overburdening at the moment. This seems to overall reflect diverse sentiments within the open source ecosystem about the regulation.

As our 2026 CRA Awareness and Readiness Report shows, many are not ready for the obligations imposed by the law, with 66% of respondents surveyed being unaware of the law, and 4 out of 10 organisations being unsure if it applies to their activities and products. However, the report notes a significant surge in published CVEs in 2026 and shows that the more diverse an open source project is, the better its security posture. Many manufacturers have been working on adjusting how they manage their cybersecurity protocols and practices in preparation for the full application of the law, such as eliminating private forks, automatic dependency updates, deploying scanning tools, and others. While the full impact of the CRA is still unfolding, 63% of respondents of our survey said that regulatory pressure changed their organisation's investment in open source dependency management (Figure 23). These are likely trends to continue with

the full application of the law.¹⁸

Overall, the number of EU legislation in the digital area is growing every year, just like the requirements to those providing services and placing products on the EU digital market. It is the same with open source, referenced in EU law with increasing frequency. Viewed as regulatory overburden by some, many consider it a natural progression and sign of maturity of tech and open source ecosystems alike. At the same time, significant efforts are underway at the EU level to harmonise policies through the Omnibus initiatives in a variety of sectors such as taxation, energy, and defense, with key digital simplification initiatives for 2026 being the AI Digital Omnibus (adopted in June 2026) and the Data Omnibus. Although simplification is widely recognised as essential for strengthening Europe's homegrown technology market, civil society and industry organisations have warned that relaxing certain safeguards may jeopardise citizens' digital rights and introduce more uncertainty.


The second half of 2026 will be busy in digital policy, with major files being proposed in the fall: the Public Procurement Act, Digital Fairness Act, European Innovation Act, as well as the overarching Multiannual Financial Framework (MFF) setting the EU's long-term budget planned to be finalised by the end of 2026. It is not yet clear how deeply these initiatives will prioritise open source. For instance, the European Sovereign Tech Fund calls for funding support for sustaining critical open source digital infrastructure through the MFF but it remains an open question whether this ambition will be backed with budget allocations.

The European Product Act (Q3 2026), comprising the revision of Regulation on European standardisation ("1025") is surely going to keep open source actors engaged. It has been long raised by open source communities that it is difficult to meaningfully

FIGURE 23

Regulation is increasing pressure to improve dependency management

To what extent has regulatory pressure (CRA or similar) changed your organisation's investment in open source dependency management?



Source: 2026 World of Open Source Survey, Q29, n = 180, Europe only, excl. 61 don't know.

¹⁸Adrienn Lawson, "2026 CRA Awareness and Readiness Report," foreword by Roman Zhukov, The Linux Foundation, June 2026.

participate in traditional standardisation systems, as pointed out in the Linux Foundation Europe response to the public consultation for this regulation at the end of 2025, a view shared by many stakeholders in the ecosystem. Open source community collaboration should be recognised as a complementary standardisation pathway. The Linux Foundation, and other open source foundations and communities, steward many de facto standards that currently cannot be referenced in the EU legislation due to the lack of formal specification, in result stifling European innovation as open source development moves way faster than complex, multi-year standardisation processes.

In our last year's report we already highlighted the need for open source experts to be more present at the policy table, and this need is only growing. The CRA did not avoid mishaps and imperfections in understanding and collaboration with the open source ecosystem, but it did provide the first serious exercise for such cooperation. Our expert interviews confirm the sentiment that the CRA has opened the dialogue with the open source community is common, and previously almost non-existent. These experiences can be leveraged by open source stakeholders into the process of negotiating and, in the future, enforcement of the Tech Sovereignty Package. The upcoming year (or three, as the typical EU legislation cycle) is going to be decisive for policy impact and leveraging of open source, if it is to be used more strategically for EU's digital sovereignty goals.

Digital policy moves towards verticals

While policy initiatives described in previous chapters are mostly horizontal, it is crucial to notice the trend of moving towards more vertically focused developments. Achieving technological autonomy requires applying open source strategies to specific economic verticals where vendor lock-in poses systemic risk. Existing frameworks such as the NIS2 Directive provide a natural foundation for identifying priority sectors such as energy, public administration, and digital infrastructure, as well as Important Entities like manufacturing and research. Across

these verticals, organisations encounter similar structural barriers in Europe: heavy reliance on a narrow set of proprietary vendors, opaque software supply chains, escalating licensing fees, and severe exit costs.

For instance, the energy sector represents a critical front for technological sovereignty in 2026, highlighted by the European Commission's Strategic Roadmap for Digitalisation and AI in the Energy Sector within the Tech Sovereignty Package. As power grids face unprecedented operational strain from integrating decentralised renewable generation, electric vehicle charging loads, and industrial electrification, traditional proprietary software procurement cycles have proven too slow to support the pace of necessary innovation. The Roadmap focuses on three pillars, including accelerating the uptake of digital and AI solutions in the EU. One of the provisions within this pillar envisages support for the development of open source AI tools to facilitate setting up single digital portals at national level to accelerate historically lengthy permit-granting procedures for energy projects, in line with the rest of the Tech Sovereignty Package.

In the area of semiconductors, the Chips Act 2.0 aims to increase the EU's share of global semiconductor production from 10% to 20% in 2030. It prioritises open silicon and the royalty-free RISC-V instruction set architecture to reduce European dependence on foreign, proprietary chip intellectual property. Implemented via the Chips Joint Undertaking and the Chips for Europe Initiative 2.0, this strategy plans to fund open-source electronic design automation (EDA) toolchains, cloud design platforms, and shared IP building blocks to empower fabless startups and domestic manufacturers. This push builds on existing commitments such as over €270 million already funneled through EuroHPC and Chips JU initiatives while directing a key portion of the broader €43 billion Chips Act public-private framework directly into open-hardware pilot lines and design infrastructure. By establishing open hardware standards across key sectors such as automotive systems, cloud infrastructure, and industrial AI, Europe aims to cultivate a self-sustaining semiconductor ecosystem where developers can customise and deploy secure microprocessors without high licensing fees or supply-chain lock-in.

Conclusions

Nearly a third of European organisations report no formal open source governance at all, and that fact runs under all three threads this report has followed: how AI is reshaping the ecosystem, how mature European open source governance actually is, and how quickly digital sovereignty has turned from talking point into mandate. These are not three separate findings but three views of the same capacity gap. Organisations that build the governance capacity to engage with open source deliberately, rather than simply consume it, capture more of what AI, contribution and procurement all have on offer. The rest are left responding to decisions made elsewhere, about their tools, their standards and increasingly their infrastructure.

AI illustrates the pattern most sharply, being the newest layer to which it has spread. Coding assistants lower the barrier to contribute and lift adoption overall, yet the same tools generate more vulnerability reports than maintainers can triage, and are compressing junior hiring even as demand for senior judgment rises. As the underlying models commoditise, the argument over openness moves up to the harness, the agents, frameworks and protocols that decide how models get used. Europe's position there echoes its position everywhere else in this report: heavy on usage and code contribution, light on governance seats, now visible at the Linux Foundation's Agentic AI Foundation. Regulation offers a partial safety net, however it is not a substitute for showing up.

The mechanism behind that gap is not a shortage of incentive. Organisations that contribute upstream report benefit-to-cost multiples of 2.3 to 3.3 times; those with advanced governance report a mean return of 4.6 times, against 3.6 times for those with none. Half the organisations that fork privately instead cannot say how many hours a release cycle costs them, so they have no figure to weigh against what contribution would have returned. Not capturing that higher return is a competitive outcome: an organisation extracting 3.6 times from the same technology

a peer extracts 4.6 times from operates at a real, compounding disadvantage, whether or not it has ever framed the gap that way. Reformed procurement, rewarding contribution over lowest cost alone, is the most direct lever available to close it at scale.

Digital sovereignty has undergone the same shift, from theoretical debate to operational mandate, over the twelve months this report covers. Ninety-four percent of European organisations call it important to their technology strategy, with security, privacy, vendor lock-in and cost risk leading as drivers, and open source is seen as central to achieving it. These findings align with landmark EU regulatory frameworks, chiefly the Tech Sovereignty Package and the Cyber Resilience Act, which require organisations to treat open technology as core strategic infrastructure. It is the first time a major policy institution has bet this boldly on open source: the Open Source Strategy's €2 billion, set against the EU's own €264 billion in annual proprietary IT spending, looks like a rounding error unless functioning OSPOs and enforced procurement rules show up alongside the money. That follow-through will be tested first in procurement: shifting from lowest-cost selection toward best price-quality and interoperability standards, and rewarding upstream contribution rather than regional origin alone, would let open solutions compete on merit and strengthen the global commons rather than wall off a regional one.

None of this argues for retreat. Europe's long-term resilience depends not on building isolated regional silos but on deepening its participation in the open source commons it already helped build, in governance seats as much as in code. The organisations, and the policies, that treat open source as infrastructure to shape are the ones this report finds capturing the return, the influence and the resilience the rest are still leaving on the table. Whether Europe closes that gap from a seat at the table, or from outside looking in, is still Europe's to decide.

Methodology


Goal	Understand the dynamics of open source across regions and industries and quantify the business value of open source software.
Data collection method	Web survey
Survey length	47 questions
Survey target population	The survey targeted employed professionals with at least a basic understanding of open source software.
Survey distribution	Linux Foundation newsletter subscribers, Linux Foundation social media channels, research partner email/social channels (Probabl, NeoNephos), panel provider
Survey fielding period	June - August 2026
Analytical sample size	367 qualified respondents for Europe and 494 for the rest of world
Margin of error (95% confidence level)	±5.1% for Europe and ±4.4% for the rest of the world
Data access	https://zenodo.org/communities/linux-foundation-research
Declaration of AI in the research process	LF Research uses generative AI and AI-assisted technologies in varying degrees for research-related tasks including aggregating secondary sources, synthesizing data, supporting report drafting, copyediting, and translation. All AI-generated content is reviewed and revised as needed.

Demographics

Figure 24 shows that respondents to the survey (sampled for Europe) worked mainly in information technology (43%), spanned a broad range of organisation sizes (38% under 250 employees, 38% between 250-4,999, 24% over 5,000) and revenue levels, and were headquartered predominantly in the EU (70%), with the rest split between the UK (21%) and non-EU Europe (10%). The sample was highly experienced with open source, where 75% reported comprehensive familiarity and 59% had over 10 years of experience. It skewed towards more technical and senior perspectives, with 88% in technical roles and 57% in leadership positions.

FIGURE 24

Respondent demographics



2026 World of Open Source Survey, Europe. Q9 n = 367; Q10 n = 365 (excl. 2); Q11 n = 280 (excl. 87); Q8, Europe n = 367, Rest of World n = 494. Regrouped. Names are shortened.

About the authors

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. His professional background encompasses entrepreneurship, executive management, software development and service as a German Air Force officer. Since joining the Linux Foundation in June 2023 as senior director for community development at Linux Foundation Europe, he drives engagement and collaboration among European open source stakeholders. He is fluent in English and German and resides in the Berlin area.

Paula Grzegorzewska is a Senior Manager, Strategic Partnerships at the Linux Foundation Europe. Her expertise revolves around digital policy and open source software, as well as digital transformation of European public services. Previously, she worked at OpenForum Europe, a leading Brussels-based think tank focused on open technologies and policy, where she was a director leading its engagement with European and global partners stemming from government, academia, not-for-profits and international institutions. She co-authored various open source policy research reports. Previously, she worked on closing the gender gap in ICT with an NGO in Luxembourg and with the start-up ecosystem in Poland, where she is from. At LF Europe, she focuses on connecting the global open source ecosystem with EU policy.

Marco Gerosa is a full professor of Computer Science at Northern Arizona University and a research analyst at LF Research. His research on software engineering and open source software has resulted in over 300 publications in top-tier venues. Dr Gerosa has a PhD, a master's in Informatics, and a B.S. in Computer Engineering. He is a Senior Member of the Institute of Electrical and Electronics Engineers (IEEE) and the Association for Computing Machinery (ACM). He supervised several PhD and MSc students who are now researchers in top institutions. He also has more than 20 years of teaching experience. For more information, visit <http://www.marcoagerosa.com>.

Acknowledgments

Authors thank the LF Research team that has provided necessary guidance, data, tools and collaboration for this report: Hilary Carter, Anna Hermansen, Adrienn Lawson, and Marco Gerosa. We also thank Thierry Carrez who leads Linux Foundation Europe for a thorough review, as well as Madalin Neag and Christopher Robinson of the OpenSSF. Thanks to Craig Ross and the Linux Foundation creative services team for their production of the PDF.

We're grateful to the stakeholders who gave their time and insights for this report: Sachiko Muto, Sebastian Wolf, Felix Reda, Volodymyr Tsap, Juan Tomás García, Lea Beiermann, Kurt Garloff, Greg KH, Leonhard Kugler, Jonas van den Bogaard. Their perspectives were invaluable in shaping our findings.

LF thanks to the following organisations who sponsored this study:



LF also thanks any survey distribution partners:



 x.com/linuxfoundation

 facebook.com/TheLinuxFoundation

 linkedin.com/company/the-linux-foundation

 youtube.com/user/TheLinuxFoundation

 github.com/linuxfoundation

October 2026



Copyright © 2026 The Linux Foundation

This report is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International Public License.

To reference this work, please cite as follows: Mirko Boehm and Paula Grzegorzewska, "2026 State of Open Source in Europe: Navigating AI and Digital Sovereignty Through Open Governance," foreword by Greg Kroah-Hartman, The Linux Foundation, October 2026.



Linux Foundation Europe is the Brussels-based European chapter of the Linux Foundation. The Linux Foundation is the world's leading home for global collaboration on open source software, hardware, standards, and data. Linux Foundation projects are critical to the world's infrastructure, including Linux, Kubernetes, Node.js, ONAP, OpenChain, OpenSSF, PyTorch, RISC-V, SPDX, Zephyr, and more. The Linux Foundation focuses on leveraging best practices and addressing the needs of contributors, users, and solution providers to create sustainable models for open collaboration. For more information, please visit us at linuxfoundation.eu and linuxfoundation.org.

The Linux Foundation has registered trademarks and uses trademarks. For a list of trademarks of The Linux Foundation, please see its trademark usage page: www.linuxfoundation.org/trademark-usage. Linux is a registered trademark of Linus Torvalds.



Founded in 2021, [Linux Foundation Research](#) explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. By leveraging project databases and networks and committing to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations worldwide.