EU Non-Proliferation and Disarmament Consortium

Promoting the European network of independent non-proliferation and disarmament think tanks

NON-PROLIFERATION AND DISARMAMENT PAPERS

No. 100 May 2026

CYBER (DIS)ARMAMENT IN PRACTICE: THE EU'S ROLE IN GOVERNING SOFTWARE VULNERABILITIES IN A FRAGMENTED INTERNATIONAL ORDER

EUGENIO BENINCASA

I. INTRODUCTION

Cyber conflict presents a profound challenge to traditional concepts of disarmament. In contrast to conventional weapons, cyber capabilities are not discrete, state-owned objects that can be counted, restricted or dismantled. They consist of evolving combinations of software, expertise, and organizational and operational practices that are largely dual-use and deeply embedded in civilian digital infrastructures. As a result, cyber capabilities resist treatment as coherent objects of disarmament in the classical arms control sense.

This structural mismatch has shifted cyber disarmament debates away from elimination and towards governance. Rather than asking how cyber capabilities can be abolished, policymakers have focused on identifying specific points within the cyber ecosystem where risks can be mitigated without undermining legitimate security, economic or innovation interests. These include technical tooling, operational infrastructure, organizational practices and norms governing use, with software vulnerabilities occupying a particularly critical position. Therefore, in this paper the term 'cyber disarmament' is not used in the narrow sense of eliminating capabilities, but as a broader analytical lens that encompasses efforts to constrain, manage and reduce the risks associated with cyber capabilities in practice, including those typically associated with arms control.

The paper argues that one of the most concrete sites where cyber disarmament can be pursued in practice is the governance of software vulnerabilities. Vulnerabilities are flaws or weaknesses in software, hardware or computer systems that attackers can exploit to gain unauthorized access, disrupt operations or cause other

SUMMARY

Cyber conflict poses fundamental challenges to traditional approaches to disarmament. Cyber capabilities are not discrete weapons, but assemblages of technical and human components, among which software vulnerabilities often serve as critical enablers of access and exploitation. This paper argues that one of the most plausible ways of pursuing cyber 'disarmament' in practice lies in the governance of software vulnerabilities, particularly through mechanisms for vulnerability disclosure. In this context, vulnerability disclosure refers to processes through which newly discovered software flaws are reported, assessed and either remediated or managed by vendors, governments and security researchers. Vulnerability disclosure does not eliminate cyber capabilities, but it shapes incentives, constrains windows of exploitation and reduces systemic risk while preserving legitimate security and innovation interests.

The paper proceeds in several steps. It first examines the structure of the global vulnerability ecosystem and the conditions that influence whether vulnerabilities are disclosed, retained or circulated. It then explains why international arms control and cyber norm processes have struggled to meaningfully engage cyber capabilities, including software vulnerabilities. Against this backdrop, the analysis shows how vulnerability governance is displaced towards domestic institutional arrangements that operate upstream of cyber operations. It turns to Europe as a case study, highlighting partial European Union-level harmonization alongside persistent national fragmentation, and concludes with recommendations for strengthening Europe's vulnerability governance framework.

ABOUT THE AUTHOR

Eugenio Benincasa is a Senior Cyber Defense Researcher at the Center for Security Studies, ETH Zürich.

adverse effects. They are not 'weapons' in themselves, but they enable many offensive cyber operations by providing access, persistence and escalation pathways across digital systems. At the same time, the accumulation and strategic retention of vulnerabilities generate systemic insecurity, exposing civilian infrastructure to exploitation by a wide range of actors, while making them especially valuable for routine intelligence collection and espionage. Their value depends on secrecy and timing, while their mitigation depends on the controlled circulation of information to those able to remediate them. As such, vulnerabilities sit at the intersection of offence and defence, public and private actors, and secrecy and transparency.

Efforts to 'disarm' from vulnerabilities do not take the form of prohibition or limitation, but of disclosure: decisions about when, how and to whom information about software flaws is revealed, including whether they are shared with vendors for patching (i.e. developing and deploying fixes to remove the flaw) or retained by governments for operational use. Vulnerability disclosure does not eliminate cyber capabilities, nor does it resolve all dimensions of cyber conflict. Even in a world with universal vulnerability disclosure, cyber operations would remain possible through social engineering, misconfigurations, supply chain compromises and other non-vulnerability-based vectors. Instead, vulnerability disclosure functions as a mechanism that shapes the incentives of states, vendors and researchers towards remediation, and reduces systemic exposure to exploitation. From a systemic perspective, disclosure is generally treated as a public good mechanism: it enables remediation and limits the accumulation of hidden risk across interconnected systems. In this sense, vulnerability governance represents a particularly concrete and policy-relevant site where cyber disarmament is pursued in practice (often unilaterally, through domestic policy choices, rather than through negotiated multilateral agreements).

This paper examines vulnerability governance as a critical—although not exhaustive—component of cyber disarmament. Section II traces the emergence of a global vulnerability ecosystem and identifies the structural conditions that shape how vulnerabilities are discovered, valued and circulated. Section III explores why international arms control and disarmament frameworks struggle to engage meaningfully with cyber capabilities, including software vulnerabilities, reflecting deeper mismatches between traditional arms control assumptions and the nature of cyber capabilities. Section IV shows how vulnerability governance is structured within domestic institutional arrangements that shape disclosure, retention and circulation decisions upstream of cyber operations. Section V maps the European Union (EU) landscape, highlighting fragmentation across legal and institutional approaches and the challenges this poses for coherent vulnerability governance. Finally, section VI derives policy recommendations for strengthening vulnerability governance in Europe.

II. THE EMERGENCE OF A GLOBAL VULNERABILITY ECOSYSTEM

In its early stages, vulnerability discovery was largely an informal activity carried out by technically skilled individuals, often motivated by curiosity, reputation or community norms.[1] Over time, however, vulnerabilities became embedded in structured economic and strategic relationships, as markets for their sale emerged, governments began acquiring them for intelligence and military purposes, and intermediaries professionalized their trade. Today, knowledge of software flaws functions as a form of capital: it can be accumulated, traded, withheld or leveraged for competitive advantage.[2] As this ecosystem matured, states increasingly moved beyond their roles as regulators or targets of cyber operations to become active participants—discovering, acquiring and retaining vulnerabilities for intelligence collection and cyber operations, while simultaneously shaping demand and exploiting opacity alongside private firms and intermediaries.[3]

To make sense of how this ecosystem operates, this paper draws on five interrelated structural conditions:

1. Talent: the availability of specialized expertise for discovering and exploiting vulnerabilities.
2. Incentives: the rewards and motivations influencing whether vulnerabilities are disclosed or retained.

[1] Benincasa, E., *From Vegas to Chengdu: Hacking Contests, Bug Bounties, and China's Offensive Cyber Ecosystem* (Center for Security Studies (CSS), ETH Zürich: Zürich, June 2024); and Perlroth, N., *This Is How They Tell Me the World Ends: The Cyberweapons Arms Race* (Bloomsbury Publishing: New York, 2021).

[2] Böhme, R., 'A comparison of market approaches to software vulnerability disclosure', Dresden University of Technology, Institute for System Architecture, 2006; and Smeets, M., *No Shortcuts: Why States Struggle to Develop a Military Cyber-force* (Hurst: London, 2022).

[3] Perlroth (note 1).

3. Circulation: the mechanisms and pathways through which vulnerability knowledge moves between actors.

4. Exclusivity: the degree to which vulnerability knowledge is restricted to a limited set of actors.

5. Transparency: the openness of information sharing and the visibility of ecosystem processes that condition how the other four conditions operate.[4]

Together, these conditions explain how the ecosystem operates and why certain outcomes—particularly limited disclosure and persistent secrecy—remain prevalent.

Talent: Human capital as the core input

The persistence of vulnerabilities is not accidental but structural. Modern digital systems are extraordinarily complex, often built from millions of lines of code, layered dependencies and legacy components, making the elimination of flaws effectively impossible.[5] However, the prevalence of vulnerabilities is not solely a function of technical complexity. The scale and persistence of software insecurity are also shaped by resource constraints across both private and public sectors, as well as commercial and regulatory environments that often reward speed, feature expansion and low cost over security, resilience and long-term maintenance. Reliance on third-party and open-source components further diffuses responsibility across supply chains, while weak liability and limited accountability have historically reduced pressure on vendors to invest in more secure development practices.[6] Even well-designed systems can generate unforeseen vulnerabilities through complex interactions between components.

These conditions ensure a continuous demand for individuals capable of identifying vulnerabilities before others do. Over time, the pool of such expertise has expanded significantly. Informal hacker communities have been complemented by university cybersecurity programmes, professional training pipelines, corporate research teams and government laboratories. Competitive formats such as capture-the-flag contests and global hacking competitions have institutionalized skill development and created visible pathways from amateur experimentation into professional security research.[7]

At the same time, advanced vulnerability discovery talent remains unevenly distributed. Certain regions, institutions and organizations concentrate disproportionate shares of high-end expertise, making human capital a strategically salient resource.[8] Because vulnerability discovery depends primarily on skills rather than physical infrastructure, it is difficult to constrain, regulate or geographically localize.

Incentives: Disclosure versus retention

Incentives ultimately determine whether vulnerabilities are disclosed, sold or retained. As noted in the discussion of talent, independent researchers play a central role in this process, as they often identify vulnerabilities that vendors and public authorities lack the capacity, visibility or incentives to discover themselves. For much of the ecosystem's history, vendors offered little or no recognition or compensation for responsible reporting. Researchers therefore faced a clear choice: disclose vulnerabilities without reward or seek buyers willing to pay for secrecy.[9] In many cases, particularly for high-impact vulnerabilities, the latter option proved more attractive.

In these segments of the market, secrecy itself became the primary source of value. For state actors, retaining exclusive access to vulnerabilities could enable long-term intelligence collection, covert surveillance or disruptive operations. The decision to retain rather than disclose typically reflects strategic calculations—the perceived intelligence or operational value of continued secrecy—rather than financial considerations. Conversely, when states do choose to disclose, this may serve aims such as signalling responsible behaviour or supporting broader stability in the digital ecosystem.[10] As a result, vulnerabilities

[4] The five conditions were identified and developed as part of a forthcoming publication by Dunn et al. (2026).

[5] Anderson, R., *Security Engineering: A Guide to Building Dependable Distributed Systems*, 3 edn (Wiley: London Dec. 2020); and Perrow, C., *Normal Accidents: Living With High-risk Technologies*, 2nd edn (Princeton University Press: Princeton, NJ, 1999).

[6] Black Duck, 'Synopsis', *2026 Open Source Security and Risk Analysis (OSSRA) Report*, 10th edn.

[7] Benincasa (note 1).

[8] Benincasa, E., *Before Vegas: The 'Red Hackers' Who Shaped China's Cyber Ecosystem* (Center for Security Studies (CSS), ETH Zürich: Zürich, July 2025).

[9] Perlroth (note 1).

[10] Ablon, L. and Bogart, A., *Zero Days, Thousands of Nights: The Life and Times of Zero-Day Vulnerabilities and Their Exploits*, RAND Research Report RR-1751 (RAND Corporation: Santa Monica, CA, 9 Mar. 2017).

with the greatest potential impact were frequently the least likely to be reported.

The expansion of bug bounty programmes and coordinated disclosure frameworks altered this incentive structure for independent researchers and security firms to some extent. By offering financial rewards, reputational benefits and legal clarity, these initiatives created legitimate pathways for monetizing vulnerability research. Some governments further supplemented monetary incentives with non-financial rewards, such as professional recognition or access to elite research environments.[11]

Nevertheless, several structural challenges persist. Vulnerability knowledge is time-sensitive, creating pressure to act quickly before value erodes through independent discovery. Transactions are shaped by information asymmetries, as buyers and sellers struggle to assess credibility without revealing too much. Pricing remains opaque, and enforcing exclusivity is difficult, particularly in informal or covert markets.[12] Ethical ambiguity also persists—once a vulnerability is transferred, its subsequent use is often beyond the discoverer's control. This loss of control can discourage disclosure to opaque buyers, but it can equally deter researchers from engaging with any intermediary—including legitimate ones—thereby reinforcing retention.

Circulation: Pathways, intermediaries and fragmentation

In the ecosystem's early years, vulnerability circulation relied on informal and often precarious channels. Underground forums, personal trust networks, closed mailing lists and private exchanges facilitated transactions, but with high levels of risk. Buyers struggled to verify authenticity without exposing themselves to fraud, while sellers faced the challenge of demonstrating value without sacrificing exclusivity.[13] With limited enforceable norms or contracts, deception and reselling were common.

As demand increased, intermediaries emerged to professionalize these exchanges. Vulnerability brokers and exploit acquisition platforms introduced pricing conventions, buyer vetting and contractual arrangements designed to reduce uncertainty.[14] While these mechanisms lowered some transaction costs, they also consolidated high-value vulnerabilities within narrow circles. Intermediaries often reinforced secrecy by channelling discoveries toward state agencies or select private clients with limited incentives to disclose them publicly.

Alongside these opaque markets, a disclosure-oriented model developed through bug bounty programmes and coordinated vulnerability disclosure (CVD) services, which facilitate responsible reporting and remediation of vulnerabilities between researchers and vendors. These mechanisms formalized responsible reporting by defining payment structures, disclosure timelines and mediation procedures between researchers and vendors. They reduced, though by no means eliminated, the legal, procedural and reputational uncertainty that had previously deterred researchers from engaging with vendors directly, creating legitimate alternatives to illicit markets. Governments also adopted similar approaches, such as establishing national vulnerability disclosure programmes and streamlined reporting mechanisms, both to strengthen defensive postures and to integrate external talent into national security ecosystems. Despite these developments, vulnerability circulation remains highly fragmented. Formal disclosure channels coexist with brokered sales and quiet retention strategies, often within the activities of the same individuals. A single researcher may responsibly disclose some vulnerabilities, sell others privately and withhold still others entirely. Circulation decisions depend on incentives, legal risk, trust and perceived strategic value rather than on any single governing norm.

Exclusivity: The strategic value of limited knowledge

In the vulnerability ecosystem, value is determined not only by the technical characteristics of a flaw, but by who knows about it and when. Exclusivity—the restriction of vulnerability knowledge to a small set of actors—is therefore central.[15] Where information itself is the commodity, scarcity increases value. A vulnerability known only to a handful of actors offers far greater operational and bargaining power than one that is widely disclosed.

[11] Benincasa (note 1).
[12] Böhme (note 2); and Smeets (note 2).
[13] Böhme (note 2).
[14] Kannan, K. and Telang, R., 'Market for software vulnerabilities? Think again', *Management Science*, vol. 51, no. 5 (2005).
[15] Kannan and Telang (note 14).

Many vulnerabilities are publicly catalogued to support tracking and remediation, notably through the Common Vulnerabilities and Exposures (CVE) system. While hundreds of thousands of vulnerabilities are documented, only a small fraction is ever observed to be exploited in practice. This gap underscores a key point: operational relevance is driven less by technical severity than by control over information.[16]

The distinction between patched vulnerabilities (N-days) and unpatched, undisclosed vulnerabilities (zero-days) illustrates this dynamic.[17] Zero-days are especially valuable because no remediation exists at the time of discovery, increasing the likelihood of successful exploitation. Their value, however, is inherently temporary and depends on continued secrecy. Once disclosed or independently discovered, their strategic utility rapidly declines.

Historically, norms of responsible disclosure sought to limit exclusivity by encouraging researchers to notify vendors privately before public release. While these norms aimed to balance user protection with researcher recognition, alternative logics gradually gained prominence. Governments, private brokers and illicit buyers increasingly treated exclusivity as an asset to be preserved rather than reduced. In these contexts, withholding information maximizes both economic returns and strategic leverage.[18]

Exclusivity thus functions as a central hinge in the ecosystem. It links talent to incentives, shapes circulation pathways and constrains the effectiveness of transparency mechanisms.

Transparency: Visibility, secrecy and strategic ambiguity

Transparency mechanisms shape which vulnerabilities become visible, how remediation is tracked and whether actors can be held accountable. Public databases, vendor advisories and coordinated disclosure processes aim to reduce systemic risk by shortening the period during which vulnerabilities remain exploitable.[19] In this sense, transparency serves a collective security function.

At the same time, transparency exists in constant tension with exclusivity. Disclosure enables patching but simultaneously destroys the operational and economic value of a vulnerability.[20] This tension produces two competing logics. A defensive logic prioritizes remediation—and the visibility needed to enable it—to reduce harm, while an offensive logic treats vulnerability knowledge as a strategic resource whose value depends on secrecy.

For many years, the extent of state participation in offensive vulnerability practices remained opaque. Public understanding shifted following major disclosures in the early 2010s, which revealed extensive retention and use of undisclosed vulnerabilities by intelligence agencies.[21] These revelations crystallized what is often described as the 'retain or disclose' dilemma: whether to preserve vulnerabilities for strategic use or disclose them to enable remediation.[22]

This dilemma highlights a structural feature of the ecosystem. States can retain vulnerabilities while avoiding public accountability for retention decisions, exploiting gaps in transparency to preserve operational flexibility. As more actors acquire comparable capabilities, transparency mechanisms themselves increasingly become objects of strategic contestation rather than neutral tools of governance.

These structural features help explain why software vulnerabilities are unusually difficult objects of security governance. Unlike many other components of cyber capability, they are latent flaws embedded in widely used digital systems and can simultaneously function as operational opportunities for some actors and systemic risks for others. They are discovered through human skill rather than industrial processes, derive value from exclusivity rather than accumulation, circulate through fragmented pathways that cut across public and private actors, and are inherently transient. Their value degrades over time as independent discovery, software updates or system changes can

[16] VulnCheck, 'State of exploitation—A peek into the last decade of vulnerability exploitation', 5 May 2024; and Cyentia Institute et al., 'A visual exploration of exploitation in the wild: The inaugural study of EPSS data and performance', July 2024.

[17] N-days are known vulnerabilities for which a patch or update already exists; zero-days are vulnerabilities that are not yet known to the affected vendor or defenders, meaning there have been zero days to develop and deploy a patch.

[18] Perlroth (note 1).

[19] Cavusoglu, H., Cavusoglu, H. and Raghunathan, S., 'Emerging issues in responsible vulnerability disclosure', Proceedings of the 18th Bled eConference, 6–8 June 2005.

[20] Smeets, M., 'The strategic promise of offensive cyber operations', *Strategic Studies Quarterly*, vol. 12, no. 3 (2018).

[21] Greenwald, G., *No Place to Hide: Edward Snowden, the NSA, and the US Surveillance State* (Metropolitan Books: New York, 2014).

[22] Leal, M. M. and Musgrave, P., 'Backwards from zero: How the US public evaluates the use of zero-day vulnerabilities in cybersecurity', *Contemporary Security Policy*, vol. 44, no. 3 (2023).

render them obsolete regardless of the holder's intent, which is arguably the most important difference to conventional arms. Their strategic utility depends less on visible deployment than on secrecy, uncertainty and the management of information asymmetries over time. As a result, the vulnerabilities with the greatest operational significance are often those least visible to external observers and least amenable to collective oversight.

These characteristics have direct implications for efforts to reduce risks associated with cyber operations. Vulnerabilities are not discrete weapons or platforms, but forms of knowledge whose value is perishable and whose control depends on practices of discovery, retention and disclosure across heterogeneous actors. Governing such an ecosystem therefore requires interventions at the level of incentives, information flows and institutional interfaces, not merely at the level of observable use.

These dynamics imply that vulnerabilities are currently governed through two partially overlapping layers. The first is market-facing and encompasses vendor obligations, CVD and disclosure pathways, including independent researchers and private security actors that shape remediation. The second is state-facing and concerns how governments decide whether to disclose or retain vulnerabilities discovered through public sector activity, or to 'disarm' from them. A broader governance gap also exists around commercial vulnerability and exploit brokers, although addressing that intermediary market lies beyond the scope of this paper. Rather, this paper focuses on the disclosure–retention nexus at the core of vulnerability governance.

These structural features explain why disclosure is the core governance problem and why multilateral 'cyber disarmament' debates have struggled to engage vulnerabilities in concrete terms. Section III shows how arms control assumptions collide with the socio-technical nature of cyber capabilities, leaving vulnerability governance largely outside the field of view.

III. THE LIMITS OF ARMS CONTROL FRAMEWORKS IN GOVERNING CYBER CAPABILITIES

If software vulnerabilities constitute a critical enabling input to contemporary cyber operations, then any credible approach to arms control must, at least implicitly, engage how vulnerabilities are discovered, retained and disclosed. Calls to regulate or limit offensive cyber capabilities have circulated for more than two decades, from early warnings of a 'Cyber 9/11' to periodic editorials advocating cold war-style limits on offensive cyber capabilities.[23] Yet international debates in cyber-focused multilateral forums have largely failed to engage with these proposals, and where they have, they have struggled to make meaningful progress. This omission is not simply an oversight or a lack of political will. It reflects deeper structural mismatches between traditional arms control frameworks and the nature of cyber capabilities more broadly.[24]

Arms control regimes were developed to regulate capabilities that are observable, attributable and comparatively stable over time. They presuppose identifiable objects of control, clear distinctions between possession and use, clear distinctions between military and civilian application, and governance mechanisms anchored in state authority, monitoring and verification. Cyber capabilities—such as the vulnerabilities that enable them—fit poorly within these assumptions.[25] They are intangible, rapidly depreciating, structurally dual-use and often embedded in private or transnational ecosystems beyond direct state control. Their strategic value depends less on accumulation than on selective non-disclosure, and less on use than on the credible possibility of use under conditions of uncertainty.[26]

Vulnerabilities make these tensions particularly visible. Although they are only one component of broader cyber capabilities, they sit at the intersection of offence and defence, public and private actors, and secrecy and transparency. As a result, they expose the limits of arms control logics more clearly than many other elements. Understanding why arms control frameworks struggle to engage cyber capabilities in general is thus a necessary step towards explaining why vulnerability governance is largely dealt with within domestic institutional arrangements.

[23] Clarke, R. A. and Knake, R. K., *Cyber War: The Next Threat to National Security and What to Do About It* (Ecco: New York, 2010); and Editorial Board, 'Arms control for a cyberage', *New York Times*, 26 Feb. 2015.

[24] Libicki, M. C., *Cyberdeterrence and Cyberwar* (RAND Corporation: Santa Monica, CA, 10 Sep. 2009).

[25] Perkovich, G. and Levite, A. E. (eds), *Understanding Cyber Conflict: 14 Analogies*, Carnegie Endowment for International Peace (Georgetown University Press: Washington, DC, 2017).

[26] Smeets (note 2).

Challenges in establishing cyber norms

The United Nations has been the primary international forum for negotiating norms governing state behaviour in cyberspace. Since 2004, a series of UN groups of governmental experts (GGE)—most recently titled the GGE on Advancing Responsible State Behavior in Cyberspace in the Context of International Security—have examined how information and communications technology (ICT) developments affect national security and military affairs. They have been given a mandate to recommend norms, rules and principles for responsible state behaviour, to examine the applicability of international law to cyberspace and to propose confidence-building measures and capacity-building initiatives.[27]

Following the 2007 cyber operations against Estonia and the 2008 conflict in Georgia, cyberspace gained recognition as a core security domain rather than a peripheral technical issue.[28] In 2013 the GGE formally affirmed that international law applies to cyberspace, and in 2015 agreed on a set of 11 voluntary, non-binding norms of responsible state behaviour—later endorsed by all UN member states through a UN General Assembly resolution—alongside confidence-building measures aimed at enhancing stability and security.[29]

Among these norms was a general commitment to responsible vulnerability handling. However, this commitment remained high level and aspirational. It did not specify disclosure timelines, clarify responsibilities between states and private actors or address the tension between vulnerability disclosure and state retention for intelligence or military purposes. As a result, vulnerabilities were acknowledged rhetorically but left largely ungoverned in operational terms.

More broadly, the implementation of cyber norms remains incipient, while the application of international law in cyberspace—although more advanced—remains contested, due to divergent conceptions of cybersecurity itself. Western states tend to frame cybersecurity around the protection of networks and information systems and defend the free flow of information as integral to the internet's architecture. By contrast, other states, including China and Russia, adopt a broader conception that explicitly includes control over information content deemed threatening to social or political stability, advancing a sovereignty-centred model of internet governance.[30] These differences complicate agreement on what constitutes offensive or defensive cyber activity and limit states' willingness to endorse norms that might constrain their freedom of action. At the same time, however, nothing prevents individual states from proactively implementing the GGE norms.[31]

These tensions culminated in the failure to reach an agreement at the fifth GGE in 2017, after China and Russia rejected references to self-defence, countermeasures and international humanitarian law (IHL), arguing that such language could legitimize the militarization of cyberspace. In response, parallel initiatives emerged. Technology companies proposed voluntary frameworks such as the Digital Geneva Convention and the Charter of Trust, while states launched the open-ended working group (OEWG) in 2018 to broaden participation beyond the limited GGE format.[32]

The OEWG process proved more inclusive and led to a consensus report in 2021 that reflected greater convergence on the applicability of IHL in cyberspace.[33] The latest OEWG mandate for the period 2021–25 established a permanent mechanism intended to provide continuity and inclusiveness in discussions on responsible state behaviour in cyberspace, which will begin its work in 2026. Nevertheless, neither process advanced concrete mechanisms for governing vulnerability discovery, retention or disclosure.

Parallel initiatives outside the UN, such as the Pall Mall Process launched by the United Kingdom and France in 2023, have focused on principles governing the development, acquisition and use of cyber intrusion capabilities, including explicit attention to vulnerabilities and exploit marketplaces. While the Pall Mall Process remains normative and non-binding, it seeks to arrive at more concrete results than previous efforts, including a code of practice for industry

[27] United Nations, General Assembly, Report of the Group of Governmental Experts on Developments in the Field of Information and Telecommunications in the Context of International Security, A/70/174, 22 July 2015.

[28] Markoff, J., 'Before the gunfire, cyberattacks', *New York Times*, 12 Aug. 2008.

[29] United Nations, General Assembly (note 27).

[30] Lieberthal, K. and Singer, P. W., *Cybersecurity and US–China Relations,* John L. Thornton China Center at Brookings (Brookings: Washington, DC, Feb. 2012).

[31] Paulus, A., 'Why Germany should practice the cyber norms it preaches: The case of vulnerabilities equities', Heinrich-Böll-Stiftung, 2022.

[32] Paulus, A., 'The global cyber norms debate', *Building Bridges in Cyber Diplomacy: How Brazil Shaped Global Cyber Norms* (Springer: Cham, 2024).

[33] Gavrilović, D., 'What's new with cybersecurity negotiations? The 2021 UN GGE report', 6 June 2021.

that is expected to address vulnerability handling alongside other issues, and a code of practice for states that includes several references to vulnerability management.[34]

Despite two decades of multilateral dialogue, clear and enforceable rules governing offensive cyber operations remain elusive. The difficulty lies not only in political disagreement, but in the mismatch between arms control assumptions and the nature of cyber capabilities themselves, particularly those rooted in software vulnerabilities.

Cyber capabilities as assemblages

A central reason why arms control frameworks struggle to engage software vulnerabilities lies not merely in how offensive cyber capabilities are defined, but in how they exist in practice. Offensive cyber capabilities are not unitary objects, but assemblages of distinct and interdependent components, including software vulnerabilities.[35]

Cyber operations depend on the interaction of 'information technology, skills, and organizations', rather than on discrete technical artefacts alone.[36] Cyber capability can be further disaggregated using the PETIO framework into people, exploits, tools, infrastructure and organization, underscoring that offensive cyber capability emerges from the interaction of multiple, often inseparable elements rather than from discrete weapons.[37] Within this framework, software vulnerabilities are best understood as an important, albeit not necessary, enabling input located primarily within the exploits component, but they do not, on their own, amount to an offensive cyber capability. Vulnerabilities acquire strategic significance only when combined with skilled operators, tailored tooling, operational infrastructure and organizational capacity capable of discovering, weaponizing, deploying and sustaining their use. From an arms control perspective, this disaggregation matters because it renders vulnerabilities legible only as one transient input among many, rather than as a discrete object of regulation. Equally, any meaningful disarmament effort would need to address all other components of the assemblage—people, tools, infrastructure and organization—not only vulnerabilities.

The people component encompasses a wide range of expertise, including vulnerability researchers, developers, operators, linguists, planners and legal advisers. The exploits component includes the code, techniques and delivery mechanisms used to leverage vulnerabilities, which may range from patched and unpatched N-days to zero-days, the latter offering particular operational advantage because they remain unknown to vendors. Tools may be custom-built for stealth or repurposed from publicly available resources. Infrastructure includes target databases and cyber ranges, while organization refers to the institutional arrangements—including rules, norms and standard operating procedures—that integrate intelligence, military and political objectives. This dispersion of capability across human, technical and organizational components further complicates any attempt to isolate vulnerabilities as governable 'weapons' within arms control frameworks.

Even recent initiatives such as the Pall Mall Process avoid defining cyber intrusion capabilities as unitary objects. Instead, they disaggregate components such as vulnerabilities and exploit development, implicitly acknowledging the difficulty of treating them as arms control objects. This component-based and deeply entangled nature of offensive cyber capabilities has direct consequences for arms control practice. Rather than governing discrete, observable objects, arms control would need to engage with capabilities that emerge from the interaction of people, knowledge, tools, infrastructure and organizational processes. The implications of this assemblage structure are not merely conceptual, they directly constrain the feasibility of applying traditional arms control mechanisms to cyberspace.

The practical limits of cyber arms control

Because offensive cyber capabilities exist as assemblages rather than discrete weapons, the core mechanisms on which traditional arms control relies encounter persistent operational obstacles. Four challenges are particularly salient, alongside the broader absence of a dedicated international legal framework for governing cyber capabilities.

First, assessing relative strength. Arms control typically relies on the ability to assess relative

[34] Paulus, A., 'Tackling the proliferation of cyber intrusion capabilities', Lawfare, 4 June 2025.

[35] Slayton, R., 'What is the cyber offense–defense balance?', *International Security*, vol. 41, no. 3 (2017); and Smeets (note 2).

[36] Slayton (note 35).

[37] Smeets (note 2).

capabilities. Because cyber capabilities take the form of assemblages rather than discrete weapons, such as nuclear warheads or conventional munitions, they cannot be counted. In some cases, certain components—such as exploit code and tooling—can be duplicated and shared at negligible cost through code reuse and digital transfer.

Second, verification of compliance. Verification would require unprecedented access to national networks, blurring the line between compliance monitoring and espionage. In an assemblage-based capability, such access would expose not only deployed tools but also underlying exploit knowledge, organizational processes and defensive weaknesses.

Third, enforcement. Enforcing compliance is equally problematic. Attribution remains contested, both technically and politically. Even when technical attribution is possible, identifying responsible human actors is far more difficult—given that assemblage-based capabilities distribute responsibility across individuals, teams, contractors and institutions—and the critical step of assigning political responsibility to a state remains contested.[38] Delays in detecting intrusions further weaken deterrence, as the deterrent effect of a response is likely to be diluted by the simple passage of time.[39]

Finally, political willingness. Even if these technical obstacles could be partially mitigated, meaningful arms control would require states to accept constraints on activities that are central to routine intelligence collection and cyberespionage. Because the same software vulnerabilities that enable disruptive cyber operations also underpin persistent access and intelligence gathering, states have strong incentives to preserve secrecy and operational flexibility rather than commit to binding limitations.

Such constraints make the emergence of a comprehensive cyber arms control agreement improbable. While future developments, such as improved attribution mechanisms, may alter this assessment, no rapid transformation in global cyber governance should be expected any time soon. As one cybersecurity expert cautions, 'a good deal depends on what one means by arms control. If the model were to be something like the treaties signed between the United States–NATO and the Soviet Union–Warsaw Pact . . . there is little basis for hope'.[40]

The structural features of the vulnerability ecosystem identified above (see section II) help explain why these limitations are particularly acute in the case of software vulnerabilities. Vulnerabilities are produced through dispersed human expertise rather than industrial processes, making them difficult to localize or constrain. Incentives within the ecosystem frequently favour selective retention over disclosure, especially where exclusivity enhances strategic or economic value. Circulation occurs through fragmented pathways that cut across public and private actors, while transparency mechanisms remain partial, uneven and contested. These characteristics undermine the core assumptions on which arms control rest: observability, attribution and the existence of stable objects of regulation.

In this sense, vulnerabilities do not merely fall outside existing arms control frameworks, they actively expose their limits. The very properties that make vulnerabilities valuable for cyber operations—secrecy, perishable knowledge and asymmetric information—also render them poor candidates for multilateral restraint. While international forums may acknowledge vulnerabilities as sources of systemic risk, they lack the institutional levers to govern how vulnerabilities are discovered, retained or disclosed in practice.

Rather than being governed through international prohibition or limitation, vulnerabilities are managed through domestic institutional arrangements. States intervene upstream of use by shaping disclosure incentives, structuring relationships with private researchers and firms, and establishing internal processes to adjudicate whether vulnerabilities discovered or acquired by government actors should be disclosed or retained.

To understand how vulnerability governance operates in practice, it is therefore necessary to move from international debates toward domestic environments in which disclosure decisions are actually made.

IV. VULNERABILITY GOVERNANCE AS DOMESTIC INSTITUTIONAL PRACTICE

Vulnerability governance is best understood as a set of domestic institutional arrangements that shape how

[38] Egloff, F. J., 'Contested public attributions of cyber incidents and the role of academia', *Contemporary Security Policy*, vol. 41, no. 1 (2020).

[39] Borghard, E. and Lonergan, S., 'Cyber operations as imperfect tools of escalation', *Strategic Studies Quarterly*, vol. 12, no. 3 (2018).

[40] Libicki (note 24).

software flaws are discovered, retained and disclosed in practice. National legal systems, administrative processes and structured relationships between public authorities and private actors play a decisive role in regulating vulnerability-related activity, intervening at the level of knowledge production, circulation and control.

In practice, states manage cyber risk not by prohibiting capabilities as such, but by structuring the conditions under which vulnerabilities are identified, reported, shared or withheld among researchers, firms and government agencies. Governance takes the form of rules, norms and procedures that determine who may discover vulnerabilities, how they may report them and under what conditions information about software flaws may circulate.[41]

Across jurisdictions, domestic vulnerability governance typically operates along two partially overlapping layers. The first layer is market-facing. It encompasses coordinated vulnerability disclosure (CVD) practices, vendor vulnerability management obligations, reporting pipelines, and the legal and procedural treatment of independent security researchers.[42] As a baseline condition, vulnerability discovery itself operates in a legally constrained space: across jurisdictions, actively probing information systems for flaws without the system owner's authorization is generally treated as unauthorized access, even when undertaken for defensive or research purposes. These arrangements therefore play a decisive role in determining whether vulnerability discovery feeds into remediation processes or is diverted into opaque markets. Where disclosure channels are predictable, legally protected and operationally responsive, vulnerabilities are more likely to be reported responsibly and patched. Where disclosure is legally risky, poorly coordinated or slow, incentives shift toward private sale, quiet retention or non-disclosure.[43]

Market-facing governance is therefore not merely technical. It performs an allocative function within the vulnerability ecosystem by shaping which actors are rewarded for disclosure, which bear legal or reputational risk and how quickly vulnerabilities transition from private knowledge to public remediation. Even where disclosure remains formally voluntary, institutional signals such as procurement expectations, enforcement discretion and public recognition play a central role in shaping researcher behaviour.

The second layer is state-facing. It concerns how governments handle vulnerabilities discovered or acquired through public sector activity, including intelligence collection, law enforcement operations, procurement and cooperation with private contractors. Here, vulnerabilities are treated not only as security risks to be mitigated, but also as potential strategic assets. Decisions about disclosure or retention involve balancing competing equities: the defensive benefits of patching and reducing systemic exposure versus the intelligence or operational value of continued secrecy.[44]

Unlike market-facing disclosure, state-facing vulnerability governance is typically opaque. Internal decision processes, oversight arrangements and outcome data are rarely public, even in democratic systems. Yet these decisions have systemic effects. Retained vulnerabilities increase exposure across shared digital infrastructures, while disclosure decisions shape vendor patching timelines and global risk distribution. As a result, state-facing governance plays a central role in determining how long high-impact vulnerabilities remain exploitable in the wild.

Crucially, these two layers are interdependent. Market-facing disclosure ecosystems shape the pool of vulnerabilities available to states, while state incentives influence the broader disclosure environment. Where states prioritize early access and control over vulnerability information, market-facing disclosure may be constrained or redirected. Where states signal a presumption toward disclosure and remediation, private reporting channels tend to strengthen. Vulnerability governance thus emerges from the interaction between public authority and private discovery rather than from either domain alone.

Because these arrangements are embedded in domestic legal and institutional contexts, vulnerability governance varies significantly across jurisdictions. Differences in criminal law, administrative capacity, regulatory philosophy and national security priorities produce distinct governance outcomes even where formal norms appear similar. Vulnerability governance is therefore best understood not as a global regime,

[41] European Union Agency for Cybersecurity (ENISA), *Coordinated Vulnerability Disclosure Policies in the EU* (ENISA: Heraklion, Apr. 2022).

[42] Householder, A. et al., *Coordinated Vulnerability Disclosure Guide for Researchers and Developers*, Software Engineering Institute (Carnegie Mellon University: Pittsburgh, PA, Aug. 2017).

[43] Ablon and Bogart (note 10).

[44] Ablon and Bogart (note 10).

but as a set of nationally specific institutional configurations operating within a shared technical ecosystem.

Where authority over vulnerability governance is concentrated within a single state, these arrangements can be internally coordinated. Where authority is distributed across multiple levels, governance outcomes depend on how competencies are divided and aligned. This distinction has particularly important implications for the EU, where authority over cybersecurity is divided rather than centralized.

V. MULTI-LEVEL AND FRAGMENTED VULNERABILITY GOVERNANCE IN THE EU

The European Union offers a particularly instructive case for examining how vulnerability governance operates in practice. As a multi-level polity, it illustrates both the possibilities and the limits of harmonizing governance across jurisdictions that retain independent authority over criminal law and intelligence activities. The analysis that follows applies the two-layer framework developed above—market-facing and state-facing governance—to trace where EU-level harmonization has advanced and where fragmentation persists.

EU-level legislation increasingly regulates product security, manufacturer obligations and institutional coordination across the single market, while decisions affecting independent vulnerability researchers and state-held vulnerabilities remain largely within national competence.[45] Recent measures, including the NIS2 directive on cybersecurity (Directive (EU) 2022/2555), the Cyber Resilience Act (CRA; Regulation (EU) 2024/2847) and the evolving Cybersecurity Act (CSA; Regulation (EU) 2019/881) framework, have strengthened EU-level coordination, reporting obligations and product security requirements. NIS2 designates national vulnerability disclosure coordinators and mandates the establishment of the EU Agency for Cybersecurity's (ENISA) European Vulnerability Database—a centralized platform for collecting, organizing and sharing information on disclosed software and hardware vulnerabilities to support coordinated response and mitigation across the EU. The CRA imposes lifecycle vulnerability management and reporting duties on manufacturers of products with digital elements. The CSA structures ENISA's mandate and the EU cybersecurity certification regime.

Yet harmonization remains limited in certain respects. While the CRA establishes significant obligations for manufacturers of products with digital elements, including mandatory notification of actively exploited vulnerabilities to ENISA and national computer security incident response teams (CSIRTs)—which handle incident monitoring and response—these instruments do not harmonize substantive criminal law, establish EU-wide safe harbour protections for independent researchers, or regulate how member states handle vulnerabilities discovered or acquired by government actors. The result is a multi-layered vulnerability governance ecosystem in which harmonization advances in product regulation and information infrastructure, while fragmentation persists in relation to researcher liability and state vulnerability practices.

Building an EU-level vulnerability governance architecture

At EU level, vulnerability governance operates through the interaction of product regulation and institutional coordination mechanisms. These instruments structure how vulnerabilities are reported, shared and remediated across the single market.

The CRA is a directly applicable regulation establishing uniform cybersecurity requirements for hardware and software products with digital elements placed on the EU market. Its core objective is to ensure that products are designed, developed and maintained with security in mind throughout their lifecycle. Central to this framework are mandatory vulnerability management and reporting obligations imposed on manufacturers. From 11 September 2026, manufacturers are required to report actively exploited vulnerabilities and severe security incidents to a national CSIRT and to ENISA via a single reporting platform. Reporting must follow tight timelines: an early warning within 24 hours of becoming aware of exploitation, a detailed report within 72 hours and a final report within 14 days of a fix becoming available.[46]

[45] Christou, G., *Cybersecurity in the European Union: Resilience and Adaptability in Governance Policy* (Palgrave Macmillan: London, 2022); and ENISA (note 41).

[46] Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements and amending Regulations (EU) no. 168/2013 and (EU) no. 2019/1020 and Directive (EU) 2020/1828

While the CRA concentrates on lifecycle security and reporting duties for manufacturers, NIS2 requires member states to designate national vulnerability disclosure coordinators and empowers national CSIRTs to facilitate coordinated disclosure processes between researchers and vendors. NIS2 also mandates the establishment of ENISA's European Vulnerability Database, intended to centralize vulnerability information at EU level.

The CRA's reporting architecture, combined with the launch of ENISA's European Vulnerability Database in 2025, creates an EU-level information infrastructure for vulnerability coordination and situational awareness.[47] Together, these mechanisms standardize how exploited vulnerabilities are identified, reported and remediated, while facilitating rapid information sharing among member states. From a systemic perspective, this framework reduces fragmentation in vendor practices, raises baseline security expectations and strengthens the EU's capacity for early warning and coordinated response to large-scale exploitation. It also limits the circulation of insecure products within the single market by embedding vulnerability management requirements into product regulation.

A key complementary development is the proposed reform of the Cybersecurity Act, unveiled by the European Commission in January 2026. The original CSA established ENISA's permanent mandate and the EU cybersecurity certification framework. The new proposal seeks to introduce risk-based limitations that could restrict certain high-risk suppliers from parts of the single market. While not directly a vulnerability disclosure instrument, the CSA framework interacts with the CRA by shaping trust and certification requirements that may influence how vulnerabilities are assessed and managed across certified products.

At the same time, this EU-level architecture does not resolve all structural tensions in the vulnerability ecosystem. In particular, the CRA's requirement for rapid reporting of unresolved vulnerabilities to public authorities has raised concerns among industry and civil society actors. In an open letter, senior figures from more than 50 organizations, including Google, Trend Micro and the Electronic Frontier Foundation, warned that this approach could create a real-time repository of unmitigated vulnerabilities accessible to multiple public authorities, potentially without adequate safeguards.[48] Experts have similarly cautioned that mandatory reporting of exploited vulnerabilities to public bodies could, under certain conditions, increase the risk of information leakage or secondary exploitation, particularly where reporting channels are insufficiently secured or sensitive information is broadly shared.

Collectively, NIS2, the CRA and the evolving CSA amount to a gradual centralization of vulnerability governance at EU level in relation to products, manufacturers and institutional coordination. Yet this harmonization remains functionally asymmetric. It standardizes reporting obligations and information flows but leaves unregulated two politically sensitive domains: the criminal law status of independent researchers and the handling of state-held vulnerabilities.

National approaches to independent vulnerability research in the EU

In contrast to the increasing harmonization of vendor obligations, the legal position of independent vulnerability researchers remains fragmented across the EU. These are individuals or teams who identify and test software flaws without a formal employment relationship or contractual mandate from the affected vendor or system owner. Where good-faith security research is not clearly protected in legislation, legality defaults to national criminal law frameworks. As a result, researcher liability remains primarily a matter of member state competence, and statutory safe harbour protections are absent in many jurisdictions.

This produces a structural asymmetry within the European vulnerability ecosystem. Manufacturers are subject to mandatory disclosure and reporting obligations, and national CSIRTs must facilitate vulnerability coordination under NIS2. Independent researchers, by contrast, often operate under general computer misuse provisions rather than research-specific legal protections.

Part of this uncertainty is rooted in the legal architecture established by the 2001 Convention on Cybercrime (Budapest Convention) and reproduced through its incorporation into national criminal law across EU member states. Article 2 criminalizes intentional unauthorized access carried out 'without

(Cyber Resilience Act), *Official Journal of the European Union*, 20 Nov. 2024.

[47] ENISA (note 41).

[48] Electronic Frontier Foundation (EFF) et al., 'Open Letter on the Cyber Resilience Act and Vulnerability Reporting', 3 Oct. 2023.

right'.[49] Although the Budapest Convention allows states to require additional elements, such as breach of security measures or dishonest intent, these limitations are optional and depend on national implementation. The framework therefore leaves the boundary between malicious intrusion and legitimate vulnerability research insufficiently defined.[50]

Directive 2013/40/EU on attacks against information systems largely mirrors the Budapest Convention's structure but introduces interpretive nuance. Paragraph 17 clarifies that criminal liability does not arise where acts lack criminal intent, including in cases of mandated security testing.[51] Nevertheless, it does not establish an explicit safe harbour for independent researchers. Protection continues to depend on domestic legislative choices and prosecutorial practice.

The EU thus leaves responsibility for managing this tension largely to member states, without imposing an obligation to introduce researcher-specific safe harbour protections. As a result, national approaches vary not only in the degree of legal protection afforded to independent researchers, but also in the availability and credibility of the disclosure pathways through which they are expected to operate. Because vulnerability governance depends on both legal and procedural arrangements, member state approaches are best understood as combinations of the two. These can be grouped into three broad policy models, as outlined below.

Policies with explicit legal provisions

A small number of states have CVD policies with explicit legal provisions to protect security researchers to different extents. Poland provides a particularly notable case. A conditional exemption from criminal liability was codified in the Criminal Code in 2017 (Article 269c), and in 2023, CSIRT NASK—the national computer security incident response team responsible for coordinating vulnerability reporting and response—together with CERT Polska (CERT.PL), its operational incident-handling unit, launched a national CVD service to operationalize disclosure procedures.[52] Article 269c exempts acts such as unauthorized access or system interference from punishment when conducted exclusively for the purpose of securing ICT systems, provided the researcher promptly notifies the system administrator and no public or private interests are harmed. France embeds vulnerability disclosure in statutory law and provides a formal legal pathway for researchers who report vulnerabilities in good faith and through the national cybersecurity agency (ANSSI).[53] Lithuania adopted a similarly formalized model in 2021 through amendments to its Cybersecurity Law, specifying in detail what constitutes lawful vulnerability research and imposing procedural obligations such as strict reporting timelines and behavioural constraints.[54]

While these frameworks appear to offer relatively high levels of legal certainty, their practical implementation warrants closer scrutiny. In the Polish case, key concepts such as 'exclusivity of purpose', 'immediate notification' and 'no harm' remain open to interpretation, potentially limiting legal certainty in practice.[55] Belgium offers a further example of this model, but one that also illustrates its practical limits. In 2023, it introduced a formal criminal law exclusion for ethical hacking through its CVD framework administered by the Centre for Cybersecurity Belgium (CCB).[56] Although the framework was intended to provide legal certainty for researchers who complied with a defined set of procedures, including a requirement that disclosure and publication be subject to prior authorization by the CCB, in practice it has struggled to do so. Critics have argued that some reported vulnerabilities were never addressed and that researchers continued to face the threat of criminal action.[57] In practice, limited institutional capacity to assess reports and authorize disclosure may undermine

[49] Council of Europe, Convention on Cybercrime (Budapest Convention), European Treaty Series no. 185, 23 Nov. 2001.

[50] ENISA (note 41).

[51] Directive 2013/40/EU of the European Parliament and of the Council of 12 August 2013 on attacks against information systems and replacing Council Framework Decision 2005/222/JHA, *Official Journal of the European Union*, L218/8, 14 Aug. 2013.

[52] CERT Polska/CSIRT NASK, 'Coordinated Vulnerability Disclosure (CVD) Policy', accessed 2 Apr. 2026.

[53] French Government, 'Loi n° 2016-1321 du 7 octobre 2016 pour une République numérique. Journal officiel de la République française' [Law no. 2016-1321 of 7 October 2016 for a digital republic, Official Journal of the French Republic], 8 Oct. 2016; and ENISA (note 41).

[54] Lithuanian Government, 'Lietuvos Respublikos kibernetinio saugumo įstatymas' [Republic of Lithuania Cybersecurity Law], as amended 28 June 2021; and ENISA (note 41).

[55] Rampášek, M. et al., 'Hunting for vulnerabilities: Call for European protection of security researchers', *Journal of Cybersecurity*, vol. 12, no. 1 (2026).

[56] Centre for Cybersecurity Belgium, 'Coordinated vulnerability disclosure (CVD)', accessed 2 Apr. 2026.

[57] De Vaere, P., 'Belgian CVD is deeply broken', 8 July 2025.

the regime's effectiveness and delay timely public disclosure.[58]

Policies without safe harbour protections

Some member states have CVD policies that do not establish statutory safe harbour protections for independent researchers. The Netherlands is emblematic of this model. Dutch CVD practice has developed largely through non-binding guidelines and institutional mediation rather than explicit legislation. National bodies such as the National Cyber Security Centre (NCSC) and the Digital Trust Center (DTC) act as intermediaries when disclosure to vendors fails, but legal protection for researchers rests on prosecutorial discretion and compliance with policy guidance rather than codified immunity. This has enabled a flexible and participatory disclosure ecosystem, but one that remains legally uncertain and dependent on trust-based arrangements.[59]

Germany follows a similar approach. Since 2022, the Federal Office for Information Security (BSI) has increasingly functioned as Germany's principal public coordination point for vulnerability disclosure.[60] Through its CVD policy, the BSI specifies expected researcher behaviour and indicates that it will not seek criminal enforcement where researchers comply with the policy's principles, except in cases involving recognizable criminal intent. Spain operates a fragmented system without a unified national CVD policy, although the two main CSIRTs (the National Cryptologic Center, CCN-CERT, and the National Cybersecurity Institute, INCIBE-CERT) maintain their own disclosure procedures. Coordination is split between CCN-CERT for the public sector and INCIBE-CERT for private companies and citizens.[61] While INCIBE-CERT outlines 'actions not allowed in the search for vulnerabilities', Spanish law does not clearly exempt good-faith unauthorized access undertaken for vulnerability research. This likely contributes to legal uncertainty for researchers and may discourage open vulnerability reporting, as no explicit safe harbour protections exist. In Romania, the National Cyber Security Directorate (DNSC) operates a national CVD programme and serves as the designated coordinator. Its vulnerability reporting guidelines include prescriptive behavioural constraints. Researchers are expected to limit their actions to demonstrating the existence of a vulnerability so as not to exceed legally permissible conduct, and to cease further testing once a flaw enabling system access has been identified rather than using it to identify additional vulnerabilities.[62]

These frameworks may improve procedural clarity and coordination, but they do not establish explicit legal protections for independent vulnerability researchers. How such assurances translate into legal certainty or prosecutorial outcomes in practice remains unclear. The German case illustrates these limits particularly well. While the BSI's policy signals institutional restraint towards good-faith disclosure, it operates at the level of administrative guidance and does not modify underlying criminal law. Researchers nevertheless remain subject to Section 202c of the Criminal Code at the time of writing, whose broad formulation can generate ambiguity around the legality of certain forms of security research. Two years after the adoption of the BSI framework, this gap continues to attract criticism from researchers and industry actors, who argue that reliance on policy-based assurances rather than statutory carve-outs continues to discourage responsible disclosure and leaves legal risk unresolved.[63]

Policies that are underdeveloped

In several member states, national CVD frameworks remain underdeveloped, partial or still in progress, leaving the status and protection of independent vulnerability researchers largely unaddressed. At the time of writing, these appear to include Italy, Ireland, Sweden and Croatia, where comprehensive statutory safe-harbour protections for independent vulnerability research remain absent and national-level policy frameworks for acceptable researcher conduct or engagement appear limited, incomplete or still in development. While some have designated national coordinators or have introduced institutional

[58] Van der Horst, M., Jansen, R. and Scherpenisse, W., 'Coordinated Vulnerability Disclosure en notificatie in het licht van NIS2' [Coordinated vulnerability disclosure and notification in light of NIS2], 27 Sep. 2025.

[59] Van der Horst, Jansen and Scherpenisse (note 58).

[60] German Federal Office for Information Security (BSI), 'Leitlinie des BSI zum Coordinated Vulnerability Disclosure (CVD)-Prozess' [BSI guideline on the coordinated vulnerability disclosure (CVD) process], 1 Dec. 2022.

[61] Spanish National Cybersecurity Institute (INCIBE), 'Vulnerability disclosure policy', INCIBE-CERT, accessed 2 Apr. 2026.

[62] Romanian National Cyber Security Directorate (DNSC), 'Coordinated vulnerability disclosure (CVD)', accessed 2 Apr. 2026.

[63] Cybersecurity Advisors Network, 'Open letter—support for responsible cybersecurity vulnerability disclosure in Germany', 28 Feb. 2026.

disclosure processes, often in the context of NIS2 transposition, vulnerability disclosure in practice remains shaped by ad hoc arrangements, sectoral initiatives or private disclosure channels. This can leave researchers exposed to varying degrees of legal and procedural uncertainty.

Government vulnerability disclosure practices across member states

EU-level measures do not meaningfully address government vulnerability disclosure practices, understood as the internal processes through which states decide whether vulnerabilities discovered or acquired by government actors are disclosed for patching or retained for intelligence, law enforcement or national security purposes. These decisions remain almost entirely within the competence of member states. Unlike the Vulnerabilities Equities Process (VEP) of the USA—an interagency mechanism through which the US government assesses whether vulnerabilities discovered or acquired by state actors should be disclosed for patching or retained for intelligence or operational use—the EU has no common framework guiding how offensive and defensive equities should be balanced, nor does EU law impose transparency or procedural requirements for state-held vulnerabilities.[64]

This absence reflects structural limits rather than oversight. The EU lacks authority over national intelligence and security services, making EU-level implementation of a government disclosure decision process impractical. As one of the developers of the VEP has noted, EU institutions such as the European Parliament and the European Commission have no mandate over member state intelligence agencies, rendering the transplantation of a US-style VEP model 'challenging'.[65] In European policy debates, analysts have expressed similar views and argued that government vulnerability disclosure decisions are better handled at the national level, given member states' direct control over their security institutions and fuller visibility into their operational landscapes.

In practice, most EU member states have not articulated explicit government vulnerability disclosure decision frameworks. The Netherlands stands out as an exception in the European context. In 2018, the Dutch government publicly acknowledged the existence of a government vulnerability disclosure decision framework, often referred to as the Dutch Vulnerability Disclosure Process or 'melden, tenzij' (disclose, unless). This framework articulates a presumption in favour of disclosure, subject to narrowly defined exceptions where retention is deemed necessary for national security or law enforcement purposes.[66] While details of individual decisions remain classified, the Dutch model stands out for its relatively high level of public articulation compared to other EU member states and for explicitly recognizing the trade-off between offensive utility and systemic security risk. This process operates entirely at the national level and is not embedded in EU law or subject to EU-level oversight.

The UK, while no longer an EU member state, provides a useful comparator given its close integration with European cybersecurity and intelligence communities. The UK has publicly acknowledged the existence of a government vulnerability disclosure framework, commonly referred to as the Equities Process, overseen by the National Cyber Security Centre (NCSC) in coordination with intelligence and law enforcement agencies.[67] Like the USA's VEP, the UK's process is designed to weigh the security benefits of disclosure against the operational value of retention. However, public information about the criteria, institutional dynamics and outcomes of the process remains limited. The UK's case, like that of the USA, nonetheless illustrates that formalized government vulnerability disclosure mechanisms can coexist with advanced offensive cyber capabilities, even if transparency remains constrained.

VI. RECOMMENDATIONS FOR STRENGTHENING VULNERABILITY GOVERNANCE IN EUROPE

If vulnerability governance is one of the most realistic levers for reducing systemic cyber risk in a fragmented international order, Europe should prioritize completing and stabilizing the institutional environment in which vulnerability discovery and

[64] White House, 'Vulnerabilities Equities Policy and Process for the United States Government', 15 Nov. 2017; and Herpig, S. and Schwartz, A., 'The future of vulnerabilities equities processes around the world', Lawfare, 4 Jan. 2019.

[65] Waterman, S., 'EU needs one set of vulnerability disclosure rules, says expert task force', Cyberscoop, 13 Mar. 2018; and Herpig and Schwartz (note 64).

[66] Dutch Government, 'Responsible disclosure', accessed 2 Apr. 2026.

[67] Levy, I., 'Equities process', National Cyber Security Centre, 5 Mar. 2025.

disclosure actually take place. The analysis above suggests that the weakest link in the European ecosystem is the uneven treatment of independent vulnerability researchers and the absence of predictable disclosure pathways across member states. Four key recommendations are outlined below.

Complete national frameworks

Europe's first priority should be the completion of national CVD frameworks in those member states where such arrangements remain partial or underdeveloped. At a minimum, these frameworks should clearly define who may report vulnerabilities, through which channels, under what behavioural expectations, and with what obligations on the receiving authority in terms of acknowledgment, triage, coordination and follow-up. Effective CVD also requires predictable engagement practices, including named points of contact, defined escalation routes, and clear expectations regarding timelines and public disclosure. Without such clarity, disclosure remains ad hoc, trust remains fragile and high-impact vulnerabilities are more likely to flow into opaque or foreign markets rather than into remediation pipelines.

Provide legal certainty

A second priority should be to provide legal certainty for good-faith security research. Across much of the EU, independent security research continues to operate in the shadow of broad computer misuse offenses, with protection resting on informal assurances rather than enforceable rights. Member states should therefore introduce explicit carve-outs or exclusions from criminal liability for good-faith security research conducted in line with recognized disclosure frameworks. Where politically and legally feasible, these protections should be codified in statute rather than policy guidance, as only statutory safeguards provide the predictability necessary to sustain long-term researcher participation. Even narrowly scoped safe-harbour provisions, conditioned on non-malicious intent, proportionality and prompt reporting, would represent a significant improvement over the current patchwork.

Build up institutional capacity

Third, and equally important in this regard, is building up institutional capacity. Designating national vulnerability disclosure coordinators and reporting channels, as required under NIS2, is insufficient if these bodies lack the resources to process reports, coordinate with vendors and authorize disclosure in a timely manner. Under-resourced coordinators risk becoming bottlenecks that delay remediation and undermine confidence in public disclosure regimes. Investment in technical expertise, staffing and procedural maturity is therefore a necessary complement to legal reform.

Encourage convergence

Finally, while the EU lacks the competence to impose a common framework for national government decisions on vulnerability retention and disclosure, it can still encourage convergence around basic principles governing how public authorities acquire, retain and handle vulnerabilities and related cyber capabilities. Greater transparency about whether member states operate with a presumption towards disclosure, even without revealing operational details, would help normalize the idea that retaining vulnerabilities entails systemic risk and therefore requires justification. The Dutch 'disclose, unless' approach illustrates that public articulation of such principles is compatible with national security concerns. While an EU-wide equities process is not feasible at present, incremental convergence at the level of principles would represent a meaningful step towards more restrained and accountable state practice.

VII. CONCLUSION

This paper has argued that software vulnerabilities are among the most concrete and policy-relevant sites where cyber 'disarmament' can be pursued in practice. Vulnerabilities are not discrete weapons that can be counted or dismantled. They are perishable forms of knowledge embedded in complex socio-technical systems, whose value is shaped by secrecy, timing and exclusivity, and whose circulation cuts across public and private actors. These properties systematically limit the usefulness of traditional arms control approaches, which depend on identifiable objects of regulation, observability and verification.

The analysis traced how the global vulnerability ecosystem is structured by five conditions—talent, incentives, circulation, exclusivity and transparency—that recurrently generate secrecy and uneven disclosure. It then showed why multilateral cyber norm processes have acknowledged vulnerabilities but have not produced operational governance mechanisms, reflecting a deeper mismatch between arms control assumptions and the assemblage-based nature of cyber capabilities. Rather than being governed through international prohibition or limitation, vulnerabilities are primarily managed through domestic institutional arrangements that shape how knowledge about software flaws is produced, retained and disclosed. In practice, this governance operates through two partially overlapping layers: market-facing disclosure and remediation systems, and state-facing decisions about whether vulnerabilities are disclosed or retained.

Europe illustrates both the possibilities and limits of this approach. EU-level legislation has significantly strengthened manufacturer obligations and institutional coordination through NIS2, the Cyber Resilience Act and the evolving Cybersecurity Act framework. Yet key segments of the ecosystem remain fragmented. The legal status of independent vulnerability researchers varies sharply across member states, and state-held vulnerability disclosure practices remain largely outside any shared European framework. In this context, Europe's most realistic pathway to reducing systemic cyber risk is not a classical arms control regime, but a more complete, predictable and trusted vulnerability governance architecture. Strengthening coordinated disclosure frameworks, clarifying researcher protections and reducing criminal-law uncertainty are therefore not peripheral technical questions, but core steps towards pursuing cyber disarmament in practice.

ABBREVIATIONS

BSI	German Federal Office for Information Security
Budapest Convention	2001 Convention on Cybercrime
CCB	Centre for Cybersecurity Belgium
CCN-CERT	Spanish National Cryptologic Center
CRA	Cyber Resilience Act
CSIRT	Computer security incident response team
CVD	Coordinated vulnerability disclosure
ENISA	EU Agency for Cybersecurity
EU	European Union
GGE	Group of governmental experts
ICT	Information and communications technology
IHL	International humanitarian law
INCIBE-CERT	National Cybersecurity Institute
NCSC	National Cyber Security Centre
N-days	Patched vulnerabilities
OEWG	Open-ended working group
VEP	Vulnerabilities Equities Process
Zero-days	Unpatched, undisclosed vulnerabilities

LIST OF RECENT NON-PROLIFERATION AND DISARMAMENT PAPERS

The Australia Group at 40: Making the AG Fit for an Era of Geopolitical Competition
Non-Proliferation and Disarmament Paper no. 99
Kolja Brockmann
June 2025

Cloud Labs and Other New Actors in the Biotechnology Ecosystem: Export Control Challenges and Good Practices in Outreach
Non-Proliferation and Disarmament Paper no. 98
Kolja Brockmann, Lauriane Héau and Giovanna Maletta
May 2025

Lessons from the EU on Confidence-building Measures Around Artificial Intelligence in the Military Domain
Non-Proliferation and Disarmament Paper no. 97
Sofia Romansky
May 2025

The EU as a Key Player in Multilateral Forums on Space Security: Perspectives for the OEWG 2025–28
Non-Proliferation and Disarmament Paper no. 96
Mathieu Bataille
April 2025

Non-proliferation, Nuclear Technology and Peaceful Uses: Examining the Role and Impact of Export Controls
Non-Proliferation and Disarmament Paper no. 95
Giovanna Maletta, Dr Mark Bromley and Kolja Brockmann
April 2025

The EU Research Security Initiative: Implications for the Application of Export Controls in Academia and Research Institutes
Non-Proliferation and Disarmament Paper no. 94
Lauriane Héau
March 2025

Subregional Arms Control and Conflict Prevention in the Western Balkans
Non-Proliferation and Disarmament Paper no. 93
Katarina Djokic
January 2025

Artificial Intelligence, Non-proliferation and Disarmament: A Compendium on the State of the Art
Non-Proliferation and Disarmament Paper no. 92
Dr Thomas Reinhold, Dr Elisabeth Hoffberger-Pippan, Dr Alexander Blanchard, Marc-Michael Blum, Dr Filippa Lentzos and Alice Saltini
January 2025

The Potentially Revolutionary Impact of Emerging and Disruptive Technologies and Strategic Conventional Weapons on the Nuclear Deterrence Debate
Non-Proliferation and Disarmament Paper no. 91
Tom Sauer
December 2024

The Nexus of Non-traditional Security and Nuclear Risk: Implications for EU Foreign Policy in the Indo–Pacific
Non-Proliferation and Disarmament Paper no. 90
Elin Bergner, Sarah Laderman and Marcy R. Fowler
November 2024

Arms Supplies to Ukraine: Does the European Arms Export Control System Need Revision?
Non-Proliferation and Disarmament Paper no. 89
Ester Sabatino
May 2024

What Happened To Demand? Getting Small Arms Control Back on Track
Non-Proliferation and Disarmament Paper no. 88
Callum Watson and Aline Shaban
March 2024

This document has been produced with the financial assistance of the EU. The contents are the sole responsibility of the EU Non-Proliferation and Disarmament Consortium and can under no circumstances be regarded as reflecting the position of the EU.

A EUROPEAN NETWORK

In July 2010 the Council of the European Union decided to support the creation of a network bringing together foreign policy institutions and research centers from across the EU to encourage political and security-related dialogue and the long-term discussion of measures to combat the proliferation of weapons of mass destruction (WMD) and their delivery systems. The Council of the European Union entrusted the technical implementation of this Decision to the EU Non-Proliferation Consortium. In 2018, in line with the recommendations formulated by the European Parliament the names and the mandate of the network and the Consortium have been adjusted to include the word 'disarmament'.

STRUCTURE

The EU Non-Proliferation and Disarmament Consortium is managed jointly by six institutes: La Fondation pour la recherche stratégique (FRS), the Peace Research Institute Frankfurt (HSFK/ PRIF), the International Affairs Institute in Rome (IAI), the International Institute for Strategic Studies (IISS–Europe), the Stockholm International Peace Research Institute (SIPRI) and the Vienna Center for Disarmament and Non-Proliferation (VCDNP). The Consortium, originally comprised of four institutes, began its work in January 2011 and forms the core of a wider network of European non-proliferation and disarmament think tanks and research centers which are closely associated with the activities of the Consortium.

MISSION

The main aim of the network of independent non-proliferation and disarmament think tanks is to encourage discussion of measures to combat the proliferation of weapons of mass destruction and their delivery systems within civil society, particularly among experts, researchers and academics in the EU and third countries. The scope of activities shall also cover issues related to conventional weapons, including small arms and light weapons (SALW).

www.nonproliferation.eu

EU Non-Proliferation and Disarmament Consortium

Promoting the European network of independent non-proliferation and disarmament think tanks

FRS Fondation pour la recherche stratégique

FOUNDATION FOR STRATEGIC RESEARCH

www.frstrategie.org

PRIF Peace Research Institute Frankfurt HSFK Hessische Stiftung Friedens- und Konfliktforschung

PEACE RESEARCH INSTITUTE FRANKFURT

www.hsfk.de

iai Istituto Affari Internazionali

INTERNATIONAL AFFAIRS INSTITUTE

www.iai.it/en

IISS

INTERNATIONAL INSTITUTE FOR STRATEGIC STUDIES

www.iiss.org/en/iiss-europe

sipri

STOCKHOLM INTERNATIONAL PEACE RESEARCH INSTITUTE

www.sipri.org

VCDNP Vienna Center for Disarmament and Non-Proliferation

VIENNA CENTER FOR DISARMAMENT AND NON-PROLIFERATION

www.vcdnp.org