American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

# Ethical Vulnerability Disclosure in Public Digital Ecosystems

**Dr. Julian R. Bennett**

Associate Professor, Australian National University (ANU), Australia

**Abstract**

Public digital ecosystems increasingly depend on interconnected government websites, identity services, cloud platforms, application programming interfaces, mobile applications, open-source libraries, software suppliers, universities, municipal systems, and digital public infrastructure. Vulnerabilities discovered within these environments can propagate beyond the organization in which they were first identified because common components and shared services create technical dependencies across institutional boundaries. Coordinated Vulnerability Disclosure provides a mechanism through which researchers, system operators, vendors, coordinators, and affected organizations can exchange vulnerability information while allowing reasonable time for remediation before detailed public disclosure. In July 2026, the U.S. Cybersecurity and Infrastructure Security Agency and international partners published updated guidance encouraging software manufacturers and online-service providers to establish formal coordinated vulnerability-disclosure programs and constructive relationships with security researchers.

This study develops an Ethical Vulnerability Disclosure Governance Framework for public digital ecosystems. The framework integrates good-faith security research, secure reporting channels, safe-harbor provisions, validation and triage, severity and exploitability assessment, multi-party coordination, remediation planning, researcher communication, public-interest disclosure, acknowledgment, and post-disclosure institutional learning. The study adopts a conceptual-methodological design supported by a transparent simulation of residual public-exposure risk across six stages of disclosure governance. These values are explicitly simulated and do not represent observed exploitation probabilities or recognized vulnerability-severity metrics. The analysis demonstrates that neither immediate unrestricted publication nor indefinite secrecy provides a universally ethical disclosure strategy. Instead, responsible disclosure requires proportionality: organizations should receive a meaningful opportunity to mitigate vulnerabilities, researchers should receive predictable protection for good-faith activity, and affected communities should ultimately obtain the information required to protect themselves. Ethical vulnerability disclosure is therefore best understood as reciprocal public-interest governance rather than a unilateral obligation imposed on security researchers.

***Keywords:*** coordinated vulnerability disclosure, vulnerability disclosure policy, cybersecurity ethics, security research, safe harbor, digital public infrastructure, vulnerability management, public-sector cybersecurity, responsible disclosure, cyber resilience

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

## 1. Introduction

Public services increasingly operate through interconnected digital ecosystems in which individual institutions rarely control every technical component on which their services depend. Government portals may rely on commercial cloud infrastructure, open-source software, external identity providers, outsourced payment services, shared authentication systems, application programming interfaces, telecommunications infrastructure, and software supplied by multiple vendors. A vulnerability in one component can therefore affect many organizations simultaneously. ENISA describes coordinated vulnerability disclosure as a process through which vulnerability finders and relevant stakeholders cooperate so that responsible parties have an opportunity to develop a fix, patch, or mitigation before broader disclosure occurs. This multi-party characteristic makes vulnerability disclosure especially important for digital public ecosystems, where one technical weakness may have consequences for public agencies, vendors, citizens, and other dependent institutions.

Security researchers are an important source of external assurance because they can identify weaknesses that internal security teams, automated scanners, procurement processes, and conventional testing have failed to detect. The practical value of their work, however, depends heavily on whether an organization has created a reliable reporting environment. The UK National Cyber Security Centre recommends that organizations establish easily discoverable and protected reporting mechanisms, publish clear vulnerability-disclosure policies, and use security.txt to advertise reporting information. RFC 9116 formalizes security.txt as a machine-parsable mechanism designed specifically to make organizational vulnerability-disclosure practices and reporting channels easier for researchers to discover. These measures appear operationally modest, yet they address a recurring governance problem: researchers cannot disclose responsibly when they cannot determine who is responsible for receiving the report.

The ethical problem becomes more difficult once a vulnerability is confirmed. Immediate unrestricted publication can alert users, researchers, competitors, journalists, and attackers at essentially the same time. In some circumstances that speed can be justified, particularly where exploitation is already occurring or an organization refuses to respond, but early publication can also increase risk when no mitigation is available. At the opposite extreme, indefinite confidentiality can permit a vulnerability to persist without accountability while users remain unaware of an exposure that affects them. CERT/CC therefore treats disclosure timing as contextual rather than fixed and allows timelines to vary according to exploitation evidence, remediation complexity, vendor coordination, and other practical circumstances. CERT/CC commonly uses a 45-day target while allowing negotiation and potentially earlier disclosure when there is significant evidence of active exploitation. Ethical disclosure is consequently a problem of balancing competing risks rather than mechanically following a universal number of days.

Researcher protection introduces another dimension. Good-faith security testing may involve probing systems in ways that superficially resemble hostile behavior even when the objective is to identify and communicate a flaw while minimizing harm. U.S. Department of Justice policy defines good-faith security research around testing, investigation, or correction conducted in a manner designed to avoid harm and used primarily to improve security or safety. European scholarship published in 2026 nevertheless argues that security researchers continue to encounter uneven legal protection across

jurisdictions, making stronger and more consistent safe-harbor arrangements an important policy issue. These conditions show why vulnerability-disclosure policy should clarify both the research behavior expected from finders and the behavior they can expect from public organizations.

## 2. Objectives of the Study

### 2.1 Establishing an Ethical Governance Model for Vulnerability Disclosure

The first objective is to formulate an ethical governance framework defining the reciprocal responsibilities of security researchers, public institutions, software vendors, vulnerability coordinators, and service providers. The model emphasizes good-faith testing, minimization of foreseeable harm, protected communication, clear scope, meaningful researcher acknowledgment, proportional remediation periods, eventual transparency, and accountability for institutional inaction.

This objective intentionally rejects the idea that responsible disclosure is solely a behavioral standard for researchers. Organizations receiving vulnerability reports also influence disclosure outcomes through the quality of their policies, technical competence, responsiveness, legal posture, and ability to coordinate remediation.

### 2.2 Strengthening Disclosure in Multi-Party Public Digital Ecosystems

The second objective is to determine how conventional vulnerability disclosure should be adapted when multiple organizations share responsibility for the affected technology.

A cloud vulnerability, shared software dependency, identity component, or open-source library can involve the original researcher, an upstream software developer, multiple downstream vendors, government operators, coordinators, and affected users. ISO/IEC TR 5895 explicitly extends vulnerability-disclosure and handling principles into multi-party settings and identifies preparation, receipt, verification, remediation development, release, and post-release activities as part of the multi-party lifecycle.

### 2.3 Evaluating Risk Reduction Through Coordinated Disclosure

The third objective is to illustrate how disclosure governance can progressively reduce residual risk to the public.

The study uses a simulated exposure-risk model rather than real vulnerability records. Its purpose is to demonstrate why reporting, triage, remediation, patch deployment, and public disclosure perform different risk-management functions and why receiving a vulnerability report should not itself be mistaken for successful vulnerability management.

## 3. Review of Literature

### 3.1 From Responsible Disclosure to Coordinated Vulnerability Governance

Vulnerability disclosure has matured substantially from informal exchanges between researchers and developers. ISO/IEC 29147:2018 provides requirements and recommendations concerning disclosure of vulnerabilities in products and services and explicitly frames vulnerability disclosure as a mechanism for reducing exploitation risk and helping users protect systems and data. ISO/IEC 30111 complements that disclosure standard by addressing internal processes for analyzing and resolving vulnerability

information. The distinction is important: one standard structures external disclosure relationships, while the other addresses the organizational handling necessary to convert reports into remediation.

CERT/CC further conceptualizes CVD as an interaction among finders, reporters, vendors, coordinators, and other affected parties. Its guidance emphasizes that the process requires more than notification; validation, communication, coordination, remediation, and eventual publication must be managed together. ENISA adopts a compatible understanding, describing coordinated disclosure as important for protecting users because vulnerabilities should normally become public after responsible parties have produced a patch, fix, or mitigating measure.

Contemporary empirical research reveals a persistent implementation gap between policy and practice. Liu and colleagues examined thousands of open-source projects and reported that although practitioner support for coordinated vulnerability disclosure was high, actual vulnerability handling frequently diverged from CVD principles, including limited availability of private reporting mechanisms. Van Hove and colleagues similarly found in 2026 that organizations with vulnerability-disclosure policies were easier to contact, yet the existence of a policy did not guarantee that reports would be resolved effectively. These studies demonstrate that disclosure policy must be backed by operational ownership and technical response capability.

The July 2020 guidance published by CISA and international partners represents a further shift toward institutionalization. The guidance encourages software manufacturers and online-service providers to establish coordinated vulnerability-disclosure programs that support productive collaboration with security researchers. This development strengthens the argument that vulnerability disclosure should be integrated into organizational cybersecurity governance rather than managed as an exceptional communications problem whenever a researcher sends an unexpected email.

### 3.2 Good-Faith Research, Safe Harbor, and Researcher–Institution Reciprocity

Good-faith security research can create public benefits by identifying vulnerabilities before they are exploited maliciously, but researchers often operate under asymmetric legal and organizational uncertainty. System owners generally know their own policies and contractual environment, whereas independent researchers may not know whether testing that is technically harmless will be interpreted as authorized research or unauthorized access.

The U.S. Department of Justice's current CFAA prosecution policy defines good-faith security research by reference to a security-improving purpose and conduct designed to avoid harm. CISA's federal vulnerability-disclosure framework likewise provides an institutional model in which defined good-faith activity can be explicitly authorized through published policy. Such mechanisms do not eliminate every legal issue, but they reduce uncertainty by replacing implicit assumptions with written expectations.

European conditions remain less uniform. Rampášek and colleagues' 2026 legal analysis argues that security researchers across Europe continue to face fragmented protection and calls for stronger safeguards supporting vulnerability reporting undertaken in the public interest. This problem is especially relevant to public institutions because aggressive responses to legitimate research may discourage future reporting and shift vulnerabilities away from defenders without actually preventing them from being discovered.

Recent empirical scholarship increasingly frames disclosure as reciprocal. Tay and colleagues' 2021 study characterizes responsible disclosure as a "two-way street," reflecting the reality that cooperation requires behavior from both researchers and vendors. Researchers should limit testing to what is necessary to demonstrate the problem, avoid unnecessary extraction of personal information, maintain confidentiality during active remediation where appropriate, and communicate sufficiently for technical reproduction. Organizations, in turn, should acknowledge reports, establish a responsible technical owner, avoid unnecessary threats against good-faith activity, provide status information, and engage honestly over remediation and disclosure.

This reciprocal approach also explains why discoverable reporting channels matter. The NCSC recommends secure web forms or dedicated email addresses and explicitly encourages publication of security.txt. RFC 9116 was created because inadequate reporting routes can cause vulnerabilities to remain unreported. The ethics of disclosure therefore begin before the researcher sends the report: organizations have a responsibility to make responsible reporting realistically possible.

### 3.3 Disclosure Timing, Multi-Party Coordination, and the Public Interest

Timing is arguably the most visible dispute in vulnerability disclosure, yet the literature suggests that it should be treated as a consequence of risk rather than as an ethical principle by itself. A fixed period can create useful expectations and prevent indefinite delay, but vulnerability complexity varies considerably. A flaw in a standalone web application may be remediated rapidly, whereas a defect embedded in widely deployed firmware, operating systems, cloud services, or supply-chain dependencies can require coordinated engineering and rollout among numerous organizations.

CERT/CC uses a target publication period while allowing modification where circumstances justify it and notes that active exploitation can support earlier disclosure. Its broader guidance identifies multi-vendor effects, response failures, and infrastructure implications as situations in which coordination can become particularly difficult. ISO/IEC TR 5895 similarly addresses multi-party CVD specifically because ordinary bilateral assumptions do not adequately describe vulnerability handling when dependent vendors and coordinators are involved.

Public disclosure should therefore be tied to public protection. Publication can create risk when exploit details appear before a defensive action is available, but it also provides important security benefits once users can patch systems, configure mitigations, assess exposure, or monitor for exploitation. ISO/IEC 29147 explicitly identifies the reduction of exploitation risk and support for user vulnerability management as purposes of disclosure.

Regulatory developments increasingly formalize the expectation that vulnerability information must move through structured channels. The EU Cyber Resilience Act establishes cybersecurity requirements for products with digital elements and introduces reporting obligations whose relevant Article 14 provisions are scheduled to begin on September 11, 2026. This does not eliminate the need for coordinated disclosure policies; rather, it strengthens the importance of organizations having processes capable of receiving, assessing, escalating, and communicating vulnerability information consistently.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

## 4. Research Methodology

### 4.1 Research Design and Evidence Selection

This study adopts a conceptual-methodological research design with simulation-based governance analysis. No live public system was scanned or tested, no non-public vulnerability report was accessed, and no unauthorized security activity was performed.

The analytical framework was constructed primarily from authoritative standards and primary guidance, including CISA's 2026 CVD guidance, the NCSC Vulnerability Disclosure Toolkit, ENISA vulnerability-disclosure resources, the CERT Guide to Coordinated Vulnerability Disclosure, RFC 9116, ISO/IEC 29147, ISO/IEC 30111, and ISO/IEC TR 5895. Peer-reviewed empirical research was used to examine operational implementation, reciprocal disclosure behavior, and researcher protection.

### 4.2 Ethical Vulnerability Disclosure Governance Framework

The proposed framework evaluates disclosure across seven linked governance components.

**Table 1: Ethical Vulnerability Disclosure Governance Framework**

| Governance Component | Primary Requirement | Organizational Responsibility | Researcher Responsibility |
|---|---|---|---|
| Good-Faith Authorization | Clear boundaries for legitimate testing | Publish scope, safe-harbor expectations, prohibited actions, and reporting conditions | Restrict testing to what is necessary and avoid unnecessary harm |
| Secure Reporting | Protected and discoverable intake mechanism | Maintain dedicated contact route, secure submission method, and security.txt | Use designated private channel where feasible |
| Validation and Triage | Rapid confirmation of legitimacy and significance | Reproduce issue, establish ownership, assess severity and exploitation evidence | Supply sufficient technical evidence for reproduction |
| Reciprocal Communication | Predictable communication during remediation | Acknowledge receipt, provide tracking information, communicate material progress | Remain reachable and protect sensitive details |
| Multi-Party Coordination | Alignment across vendors and dependent systems | Identify upstream/downstream dependencies and use coordinators where needed | Disclose affected dependencies known to the researcher |
| Remediation and Release | Effective correction before avoidable exposure | Develop, test, deploy, and communicate mitigations | Support validation or retesting where appropriate |

**Source:** Author-developed synthesis based on current CISA, NCSC, CERT/CC, ENISA, IETF, and ISO guidance.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

### 4.3 Simulation Procedure

A conceptual Residual Public-Exposure Risk Index was developed for methodological illustration. The index ranges from 0 to 100, with higher values representing greater unresolved public exposure.

The simulation assigns a risk value of 92 at initial discovery, 82 after protected private reporting, 67 after validation and triage, 45 when a mitigation is ready, 23 after patch or compensating-control deployment, and 12 after coordinated disclosure.

These values do not represent CVSS, EPSS, CISA Known Exploited Vulnerability assessments, exploit probabilities, or empirical incident rates. They are intentionally synthetic.

**The conceptual relationship can be expressed as:**

$R_{t+1}=R_t-\Delta G_t$

where $R_t$ represents residual exposure risk at disclosure stage t, and $\Delta G_t$ represents the simulated risk reduction produced by the governance intervention implemented during that stage.

## 5. Analysis

### 5.1 Comparing Alternative Disclosure Strategies

The ethical strengths and weaknesses of vulnerability disclosure become clearer when coordinated disclosure is compared with two extreme alternatives.

**Table 2: Comparative Assessment of Vulnerability Disclosure Strategies**

| Disclosure Strategy | Transparency | Pre-Patch User Protection | Risk of Organizational Delay | Researcher Predictability | Overall Ethical Assessment |
|---|---|---|---|---|---|
| Immediate Unrestricted Publication | Very High | Low where no mitigation exists | Very Low | Moderate | Justifiable in exceptional high-urgency cases but risky as a universal default |
| Indefinite Private Disclosure | Very Low | Potentially High initially | Very High | Low | Unsustainable when remediation stalls or affected users require warning |
| Coordinated Vulnerability Disclosure | Progressive | High when coordination functions properly | Moderate and manageable through agreed escalation | High where policy and safe harbor are explicit | Strongest default model for good-faith vulnerability governance |

Immediate public disclosure can be ethically defensible where active exploitation creates urgent public need for information or where responsible parties are demonstrably unwilling to act. However, it may expose users unnecessarily when a correction could have been developed rapidly through private coordination.

Indefinite confidentiality creates the opposite problem. It may initially minimize publication-related exploitation but can ultimately protect organizational convenience rather than citizens if no remediation occurs.

Coordinated disclosure offers a structured middle position. It preserves temporary confidentiality because that confidentiality has a specific protective purpose: enabling remediation. Once that purpose has been fulfilled—or when continued confidentiality no longer serves it—public disclosure becomes appropriate.

### 5.2 Funnel Analysis of Residual Public Risk

**Graph 1: Simulated Funnel of Residual Public-Exposure Risk Across Ethical Vulnerability Disclosure**

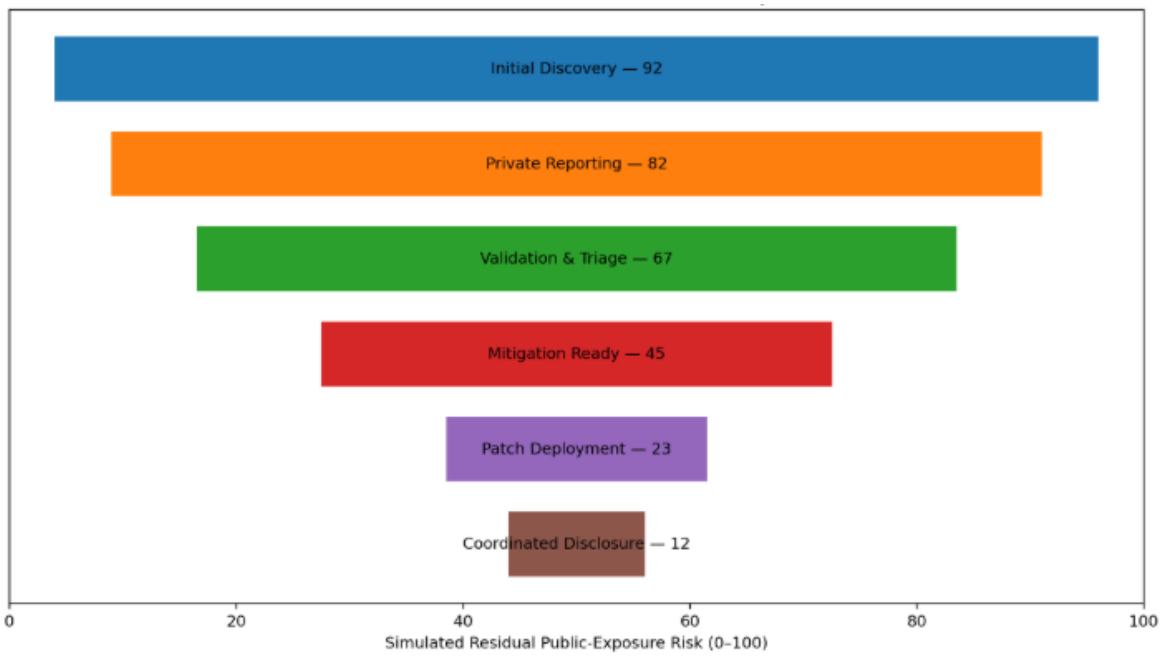


**Source:** Author-generated simulation; values are hypothetical methodological data.

The funnel begins with a simulated exposure score of 92 at discovery. Private reporting reduces the score to 82 because the issue has entered a managed process, but the underlying vulnerability still exists. This distinction is critical. An organization should not count vulnerability receipt as vulnerability remediation.

Validation and triage reduce the illustrative score to 67 because organizations can now determine the affected assets, technical cause, plausible exploitability, dependent systems, and required owners.

CERT/CC treats validation as a distinct phase because reports must be assessed for scope and credibility before remediation can proceed effectively.

The largest simulated reduction occurs as mitigation becomes technically available and is deployed. This reflects the principle that the strongest protection comes from reducing exploitable exposure rather than suppressing information.

The final score of 12 follows coordinated disclosure. The remaining value represents residual risk such as delayed patch adoption, unmaintained systems, unsupported products, or users who do not apply available mitigation. Public disclosure can reduce this residual risk by making the vulnerability, affected versions, and recommended protective actions visible.

### 5.3 Escalation, Deadlock, and Failed Coordination

Ethical disclosure frameworks must specify what happens when cooperation breaks down. Publishing a policy does not guarantee response capacity. Van Hove and colleagues' 2026 research found that policies improved contactability but did not necessarily ensure effective remediation. This reinforces the need for escalation mechanisms.

When an organization appears unresponsive, the first step should normally be to verify that the report reached an appropriate technical owner and to use alternate documented channels. Where the vulnerability affects multiple vendors or a public institution lacks coordination capacity, a national CSIRT, CERT/CC, sector coordinator, or another established vulnerability coordinator can help align stakeholders. CERT/CC specifically provides guidance for circumstances involving coordination difficulties and multi-party vulnerabilities.

If communication remains ineffective, disclosure timing should be reassessed according to actual risk. Evidence of active exploitation, broad exposure, public-safety consequences, independent rediscovery, information leakage, or lack of genuine remediation progress can strengthen the argument for accelerated disclosure.Conversely, a short extension may be ethically appropriate where remediation is demonstrably close to release and a brief delay materially reduces user harm.

## 6. Discussion

### 6.1 Ethical Vulnerability Disclosure Depends on Reciprocal Institutional Conduct

The analysis demonstrates that coordinated vulnerability disclosure should not be framed as a rule requiring researchers to behave responsibly while leaving organizations free from equivalent obligations. The effectiveness of disclosure is produced jointly. A researcher can minimize testing, avoid accessing unnecessary data, submit privately, and allow a reasonable remediation opportunity, yet the process will still fail if the organization ignores the report, sends it to the wrong department, threatens legal action reflexively, or lacks technical ownership. Conversely, even the strongest disclosure program cannot function when a reporter intentionally creates harm, withholds essential technical information, or uses vulnerability possession primarily as leverage for extortion. The U.S. Justice Department's good-faith research formulation reflects this distinction by linking legitimate research to a security-improving purpose and conduct designed to avoid harm.

Reciprocity also has practical cybersecurity value. Researchers frequently possess information that an internal security team does not have, including a reproducible trigger, proof-of-concept sequence,

environmental assumptions, or insight into how separate components combine into an exploitable state. Organizations possess complementary information regarding architecture, downstream dependencies, customer deployment patterns, and remediation constraints. Cooperation therefore improves the quality of both vulnerability analysis and remediation.

Recent research characterizing responsible disclosure as a two-way relationship is especially relevant to public digital ecosystems, where security researchers and public organizations ultimately share an interest in protecting users. A mature disclosure program should therefore measure not only time to patch but also time to acknowledgment, communication continuity, researcher experience, percentage of reports successfully triaged, dependent vendors contacted, and effectiveness of post-release mitigation.

### 6.2 Safe Harbor and Discoverable Reporting Channels Are Components of Security Architecture

Vulnerability-disclosure policy is sometimes treated as administrative documentation external to cybersecurity architecture. The evidence examined in this study supports the opposite interpretation. A discoverable vulnerability-reporting mechanism is itself a defensive control because it increases the probability that information about an unknown weakness reaches the people capable of correcting it.

RFC 9116 addresses precisely this discovery problem by standardizing security.txt. The NCSC similarly recommends combining an accessible reporting channel, a clear policy, and security.txt as core elements of a vulnerability-disclosure process. These mechanisms require little infrastructure compared with many security controls, yet their effectiveness depends on what happens behind the contact address. Reports must enter a real triage workflow rather than disappearing inside general customer support.

Safe-harbor language serves a comparable architectural function because fear of legal retaliation can distort the disclosure pathway. If researchers believe that reporting creates greater personal risk than remaining silent, vulnerabilities may remain undisclosed or may be routed through less constructive channels. European legal scholarship indicates that the protection landscape remains uneven, demonstrating that organizational clarity continues to matter even as broader regulation evolves.

Safe harbor should nevertheless be written carefully. It should not provide unlimited authorization for destructive testing. A well-designed policy should define scope, discourage persistence and service disruption, limit access to personal data, explain what researchers should do if sensitive information is encountered, and state the conditions under which the organization will regard the activity as authorized good-faith research.

### 6.3 Public-Interest Transparency Requires Neither Premature Exposure nor Permanent Secrecy

The strongest ethical position emerging from the analysis is sequenced transparency. Temporary confidentiality is justified when it materially contributes to safer remediation. Public disclosure is justified when it enables users, downstream operators, administrators, and defenders to take informed protective action. The ethical purpose of coordination is therefore to place these activities in the correct sequence rather than elevate either secrecy or publication into an absolute principle.

ENISA explicitly connects coordinated disclosure with allowing responsible parties to develop patches or mitigating measures before public release. ISO/IEC 29147 similarly explains vulnerability

disclosure in terms of reducing exploitation risk and enabling users to improve technical vulnerability management. These positions support disclosure that is timed around protection rather than reputation.

This approach becomes particularly important in public digital ecosystems because users may not have a realistic alternative to the affected service. Citizens cannot always abandon an identity portal, taxation platform, municipal service, public-health system, or educational infrastructure simply because they disagree with the operator's cybersecurity practices. Public institutions therefore have heightened responsibilities to communicate material security information where individuals need that information to reduce risk.

The approaching implementation of Cyber Resilience Act vulnerability-reporting obligations in September 2020 provides further evidence that vulnerability communication is becoming an institutional responsibility rather than solely an informal norm within security communities. Public organizations should prepare for this environment by building disclosure programs capable of handling both researcher reports and formal regulatory reporting without confusing the two processes.

## 7. Conclusion

Ethical vulnerability disclosure is increasingly important to the resilience of public digital ecosystems because contemporary public services depend on complex technical relationships that no individual organization can inspect completely. Security researchers can provide valuable independent evidence about weaknesses that internal teams, vendors, and conventional assurance processes have missed. The public value of that research, however, depends upon the existence of credible institutional pathways through which vulnerabilities can be reported and corrected before they create unnecessary harm.

The study demonstrates that an ethical disclosure model must be reciprocal. Researchers should operate in good faith, minimize unnecessary access and disruption, provide reproducible technical evidence, and preserve appropriate confidentiality while active remediation is proceeding. Public institutions and vendors should provide discoverable secure reporting channels, clear scope and safe-harbor language, meaningful acknowledgment, technically competent triage, predictable communication, multi-party coordination, and proportionate remediation. An organization that demands responsible behavior from researchers while providing no reliable means of reporting cannot reasonably claim to possess a mature vulnerability-disclosure program.

The simulation further illustrates why the disclosure lifecycle must be understood as progressive risk reduction. Residual simulated public exposure begins at 92 when the weakness remains unmitigated and declines through private reporting, validation, remediation development, patch deployment, and eventual coordinated disclosure to 12. These values are illustrative rather than empirical, but the pattern emphasizes a crucial distinction: information handling reduces uncertainty, whereas remediation reduces exploitability. Public disclosure contributes most effectively after defensive action is available, although active exploitation, institutional inaction, or other extraordinary circumstances may justify an accelerated timeline.

The broader conclusion is that ethical vulnerability disclosure is neither secrecy nor publication; it is accountable coordination directed toward reducing public harm. Future research should empirically compare public-sector disclosure programs across countries, measure response and remediation times, examine the effect of explicit safe-harbor language on researcher participation, and evaluate disclosure

quality in cloud and multi-vendor supply chains. As formal vulnerability-reporting obligations expand, institutions should use the opportunity not simply to create compliance procedures but to develop durable relationships with the security-research community. A strong public digital ecosystem is one in which responsible discovery is safe to report, credible findings are rapidly actionable, remediation is independently verifiable, and eventual transparency gives users the knowledge required to protect themselves.

**References**

1. Arora, A., Telang, R., & Xu, H. (2008). Optimal policy for software vulnerability disclosure. *Management Science, 54*(4), 642–654. https://doi.org/10.1287/mnsc.1070.0770
2. Arora, A., Krishnan, R., Telang, R., & Yang, Y. (2010). An empirical analysis of software vendors' patch release behavior: Impact of vulnerability disclosure. *Information Systems Research, 21*(1), 115–132. https://doi.org/10.1287/isre.1080.0216
3. Cavusoglu, H., Cavusoglu, H., & Zhang, J. (2008). Security patch management: Share the burden or share the damage? *Management Science, 54*(4), 657–670. https://doi.org/10.1287/mnsc.1070.0770
4. Frei, S., May, M., Fiedler, U., & Plattner, B. (2006). Large-scale vulnerability analysis. *Proceedings of the 2006 SIGCOMM Workshop on Large-Scale Attack Defense*, 131–138. https://doi.org/10.1145/1162666.1162668
5. Ozment, A. (2004). Bug auctions: Vulnerability markets reconsidered. *Proceedings of the 3rd Workshop on Economics and Information Security*, 1–16.
6. Schechter, S. E. (2002). Computer security strength & risk: A quantitative approach. *Harvard University Computer Science Technical Report*.
7. Cogo, V., & Tonella, P. (2020). A survey on vulnerability disclosure programs and bug bounty platforms. *ACM Computing Surveys*.
8. Ruohonen, J., & Allodi, L. (2021). A review of vulnerability disclosure programs and bug bounty platforms. *Computers & Security, 104*, 102247.
9. Finifter, M., Akhawe, D., & Wagner, D. (2013). An empirical study of vulnerability rewards programs. *22nd USENIX Security Symposium*, 273–288.
10. Zhao, M., Grossklags, J., & Li, P. (2014). The economics of vulnerability disclosure: Incentives and market structures. *Proceedings of the 13th Workshop on the Economics of Information Security*, 1–20.
11. Maillart, T., & Stauffer, C. (2018). A quantitative study of vulnerability disclosure programs. *Proceedings of the 17th Workshop on the Economics of Information Security*, 1–18.
12. Householder, A. D., Wassermann, G., Manion, A., & King, C. (2017). *The CERT Guide to Coordinated Vulnerability Disclosure*. Carnegie Mellon University Software Engineering Institute.
13. ISO/IEC. (2014). *ISO/IEC 29147:2014 Information Technology—Security Techniques—Vulnerability Disclosure*. International Organization for Standardization.
14. ISO/IEC. (2018). *ISO/IEC 30111:2019 Information Technology—Security Techniques—Vulnerability Handling Processes*. International Organization for Standardization.
15. Mozilla. (2019). *Security Bug Bounty Program and Vulnerability Disclosure Guidelines*. Mozilla Foundation.

16. National Institute of Standards and Technology. (2020). *Vulnerability Disclosure Framework: NISTIR 8259A*. U.S. Department of Commerce.
17. National Telecommunications and Information Administration. (2016). *Vulnerability Disclosure Attitudes and Practices Survey*. U.S. Department of Commerce.
18. ENISA. (2018). *Coordinated Vulnerability Disclosure: Benefits and Challenges*. European Union Agency for Cybersecurity.
19. HackerOne. (2020). *The Hacker-Powered Security Report 2020*. HackerOne.
20. van't Hof, C. (2016). *Fearless Digital: A Guide to Ethical Hacking and Responsible Disclosure*. Amsterdam University Press.