



A Qualitative Document Analysis of Public Vulnerability Disclosure and Bug Bounty Policies from Selected Technology Companies

MOHD RUHAIFI ZAINOL and MOHAMAD FADLI BIN ZOLKIPLI

School of Computing, College of Art and Science, Universiti Utara Malaysia (UUM), 06010 Sintok, Kedah, MALAYSIA
Email: ruhaifi.zainol@aiu.edu.my m.fadli.zolkipli@uum.edu.my | Tel: +60135932033 |+60177247779 |

Received: May 09, 2026
Accepted: May 14, 2026
Online Published: June 04, 2026

Abstract

Public vulnerability disclosure policies and bug bounty programmes are now common instruments of cybersecurity governance, but the quality of these documents varies in ways that affect researcher participation, organisational accountability, and user protection. This paper re-examines public vulnerability disclosure and bug bounty policies from selected technology companies through qualitative document analysis. The analysis treats policy documents not as routine support pages, but as organisational texts that define authorised conduct, allocate risk, communicate trust, and structure the practical relationship between companies and external security researchers. The study focuses on six dimensions: scope, reporting procedure, safe-harbour protection, reward logic, disclosure timing, and response transparency. The paper finds that public policies commonly perform two roles at once. On the surface, they invite cooperation and signal openness to external security reporting. At a deeper level, they preserve organisational control through eligibility rules, exclusions, conditional legal protection, discretionary reward decisions, and open-ended remediation timelines. This tension is consistent with prior research showing that vulnerability disclosure is shaped not only by technical discovery, but also by incentives, market design, vendor response behaviour, legal uncertainty, and trust between organisations and researchers. The paper argues that policy maturity should therefore be assessed by documentary quality rather than by visibility alone. A mature policy should provide clear scope, credible safe harbour, usable reporting instructions, transparent triage expectations, proportionate disclosure rules, and fair reward criteria. These features do not remove company control, but they make that control more predictable and accountable. The analysis concludes that disclosure and bounty policies are operational governance instruments, and that weak drafting can reduce the practical value of otherwise well-intentioned cybersecurity programmes.

Keywords: bug bounty, coordinated vulnerability disclosure, cybersecurity governance, document analysis, safe harbour, vulnerability disclosure

1. Introduction

The expansion of digital services has made it increasingly unrealistic for technology companies to rely only on internal security teams. Modern platforms change rapidly, attack surfaces span web, mobile, cloud, application programming interface, and software supply-chain environments, and external researchers may discover weaknesses before internal teams do. Vulnerability disclosure has therefore become a governance problem as much as a technical process. Earlier research on vulnerability life cycles shows that discovery, reporting, patching, and disclosure are connected events rather than isolated technical moments (Shahzad et al., 2012). This means that the value of vulnerability discovery depends heavily on how organisations receive reports, assess risk, communicate with researchers, and coordinate remediation. Public vulnerability disclosure policies sit at the centre of this relationship. They tell researchers which systems may be tested, what behaviour is permitted, where to submit findings, what information to include, and whether the organisation will protect good-faith research. Where bug bounty programmes are added, the policy also converts selected vulnerability findings into compensable work. Security economics research has shown that vulnerability information can have commercial value, but that the conditions under which vulnerabilities are traded, disclosed, or rewarded are shaped by incentives and institutional constraints (Allodi, 2017, Algarni & Malaiya, 2014). As a result, policy documents are not neutral administrative pages. They are governance instruments that manage the boundary between external security labour and organisational control.

This boundary deserves careful examination because a policy may appear welcoming while still being narrow, legally cautious, or highly discretionary. A company may invite researchers to submit reports but reserve broad authority to decide whether the testing was authorised, whether the vulnerability is eligible, whether a reward will be paid, and when public discussion may occur. The economics of vulnerability disclosure literature has long recognised that vendor decisions and researcher incentives affect whether disclosure improves security outcomes (Savage et al., 2015).



Similarly, research on patch-release behaviour shows that organisational response is a measurable factor in the impact of disclosure (Arora et al., 2010). These findings suggest that policy wording matters because it shapes the expectations that exist before any report is submitted. This paper addresses the topic through qualitative document analysis of selected public vulnerability disclosure and bug bounty policy documents from technology companies. The aim is not to rank companies or to identify one best programme. Instead, the paper analyses recurring patterns in policy design and considers how those patterns affect trust, accountability, and participation. The study is guided by three research questions. First, what core dimensions recur in public vulnerability disclosure and bug bounty policies? Second, how do selected technology companies frame the relationship between company authority and researcher cooperation? Third, what characteristics make a disclosure or bounty policy more credible as a cybersecurity governance instrument? The paper argues that the existence of a policy is not sufficient evidence of maturity. A publicly visible disclosure page may indicate awareness, but maturity depends on the quality of the commitments made in the document. The National Institute of Standards and Technology's guidance on vulnerability disclosure programmes emphasises that organisations need structured intake, triage, coordination, remediation, and communication practices rather than only a reporting address (NIST, 2023). The NIST Cybersecurity Framework 2.0 also places governance at the centre of cybersecurity risk management, which supports treating disclosure policies as part of organisational governance rather than as purely technical support material (NIST, 2024, Govern Function). From this perspective, a well-designed policy should reduce uncertainty for both the company and the researcher while preserving operational safety.

2. Literature Review

2.1 Vulnerability disclosure as coordinated governance

Vulnerability disclosure was historically discussed through the contrast between full disclosure and responsible disclosure, but contemporary practice is better described as coordinated vulnerability disclosure. Coordination recognises that immediate public release can create user risk, while indefinite secrecy can leave vulnerabilities unresolved and reduce accountability. Formal models of optimal vulnerability disclosure show that timing decisions involve trade-offs among vendor incentives, attacker awareness, and social welfare (Arora et al., 2008). Later legal and policy analysis similarly emphasises that the timing of disclosure has value because it affects both remediation incentives and public risk (Kesan et al., 2019). This governance view is important because disclosure is never only an information transfer. A vulnerability report must pass through social, legal, and organisational processes before it produces security improvement. If the reporting channel is unclear, the receiving organisation is slow, or the researcher fears retaliation, the technical validity of the vulnerability may not translate into effective remediation. Empirical work on large-scale web vulnerability notification demonstrates that even when valid information is provided, notification and remediation are difficult to scale because recipients vary in capacity, attention, and responsiveness (Stock et al., 2016). This supports the argument that disclosure policies should be evaluated by the practical processes they create, not merely by their existence.

Coordinated disclosure also distributes authority unevenly. Organisations own or control the systems, define acceptable testing boundaries, classify severity, and decide whether remediation has occurred. Researchers operate with less information and greater personal risk. Comparative analysis of coordinated vulnerability disclosure policies shows that real-world policies differ significantly in their treatment of scope, reporting expectations, disclosure timing, and legal reassurance (Disso et al., 2023). These differences matter because external researchers often use the public policy as their first evidence of whether a company is prepared to handle reports fairly.

2.2 Bug bounty programmes and vulnerability markets

Bug bounty programmes introduce an explicit incentive layer into disclosure practice. They offer rewards for valid findings and allow companies to draw on a distributed pool of security researchers. Early empirical work on vulnerability reward programmes found that such programmes can produce meaningful security reports, but their effectiveness depends on programme design, reward levels, and administrative capacity (Finifter et al., 2013). Subsequent research revisiting the idea that "given enough eyeballs, all bugs are shallow" showed that bug bounty programmes can broaden discovery, but that researcher participation and vulnerability yield are shaped by platform rules and programme maturity (Maillart et al., 2017). The market framing of bug bounty programmes is useful but incomplete. Vulnerability discovery is uncertain labour: researchers invest time without knowing whether the issue is in scope, already known, duplicated, or eligible for reward. Studies of vulnerability markets suggest that buyers, sellers, and intermediaries operate under information asymmetry, which complicates the idea of a simple open market for bugs (Kannan et al., 2016). Work on the industrialisation of vulnerability discovery also shows that vulnerability research has become more organised, professionalised, and economically structured over time (Householder & Manion, 2011). Public bounty policies therefore need to be read as market-design documents as well as security documents.

Reward decisions are especially sensitive because they convert technical findings into economic value. Severity scores, exploitability judgments, duplicate handling, and report quality requirements all affect whether a researcher is paid



(Munaiah and Meneely, 2016) argued that there can be a disconnect between vulnerability severity scoring and bounty decisions, which matters because researchers may interpret rewards as signals of value and fairness. Likewise, Laszka et al. (2018) showed that bug bounty platforms must manage misaligned incentives when validating reports. These studies suggest that transparent reward logic is not merely a researcher preference; it is part of the institutional design of a credible programme.

2.3 Safe harbour, legal uncertainty, and trust

Legal uncertainty remains a central barrier to good-faith security research. A company may request reports while researchers still worry that ordinary exploratory testing could be interpreted as unauthorised access, violation of terms of service, or misuse of data. Safe-harbour clauses attempt to reduce this fear by stating that the organisation will not pursue legal action against researchers who follow the policy and act responsibly. Research on safe harbour in cybersecurity disclosure practice frames this protection as a governance mechanism because it clarifies the conditions under which external research is tolerated or encouraged (Zolkipli et al., 2022). However, safe harbour is only credible when it is specific enough to guide behaviour. If a policy uses appreciative language but reserves broad discretion to redefine acceptable behaviour after the fact, the protection may feel symbolic rather than practical. Recent work on organisational responses to coordinated vulnerability disclosure shows that trust and control are closely linked: organisations need enough control to protect systems and users, while researchers need enough trust to believe that responsible reporting will not expose them to unfair treatment (O'Hare et al., 2024). This is why safe harbour cannot be evaluated as one isolated sentence. It must be read together with scope, prohibited conduct, data-handling rules, and response commitments.

Trust is also shaped by the information environment around vulnerability reporting. Technical threat intelligence research shows that security decision-making depends on timely, credible, and actionable information (Tounsi & Rais, 2018). A vulnerability report is one form of such information, but the willingness to provide it depends on whether the researcher believes the recipient will respond constructively. A disclosure policy that clearly explains what researchers may do, how they should report, and what protection they receive improves the conditions under which actionable information can flow.

2.4 Response practice and accountability

The receiving organisation's response behaviour is a major determinant of disclosure success. Vendor patching studies show that disclosure pressure can influence patch-release behaviour, but the effect depends on organisational incentives and capacity (Arora et al., 2010). Broader empirical work on vulnerability discovery ecosystems also shows that disclosure involves multiple actors whose interactions shape whether vulnerabilities are found, reported, validated, and fixed (Zhao et al., 2015). These findings imply that a policy should not only impose obligations on researchers; it should also state what the organisation will do after receiving a report. Many public policies are detailed about researcher restrictions but vague about company responsibilities. They may prohibit denial-of-service testing, data exfiltration, social engineering, privacy intrusion, and public disclosure, but give only general statements about reviewing reports as soon as possible. This asymmetry matters because a credible policy should distribute procedural expectations on both sides, especially because vulnerability disclosure economics depends on incentives and response behaviour (Walshe & Simpson, 2020). If researchers must invest time and accept risk while the organisation offers no response timeline or escalation path, the relationship becomes structurally one-sided.

The need for accountability is also visible in software ecosystems where vulnerabilities spread across packages, dependencies, and shared infrastructure. Ruohonen's (2021) analysis of vulnerabilities in Python packages shows how software risks often exist within broader dependency environments rather than in isolated products. For technology companies, this means disclosure policies must address not only first-party assets but also third-party services, open-source components, and supply-chain boundaries. A mature policy should explain how researchers should handle boundary cases instead of leaving them to guess whether a finding is reportable.

3. Methodology

This study applies qualitative document analysis to public vulnerability disclosure and bug bounty policy documents from selected technology companies. The method is appropriate because the research problem concerns wording, framing, institutional logic, and the way policy texts structure relationships between companies and external researchers. Document analysis allows the policies to be treated as primary governance texts. It does not assume that the public document fully represents internal practice, but it recognises that the public document strongly shapes first contact, researcher expectations, and the legitimacy of the reporting process. Document selection followed three practical criteria. First, the policy had to be publicly accessible without private onboarding. Second, the company had to operate in a technology-intensive environment where external vulnerability reporting is plausible. Third, the document had to contain enough policy substance to support thematic reading, including at least some discussion of scope, reporting method, researcher conduct, reward eligibility, or disclosure timing. These criteria are consistent with



the idea that vulnerability disclosure programmes should be assessed as structured processes rather than as isolated contact mechanisms (NIST, 2023). The analysis used six thematic dimensions: scope, reporting channel, safe harbour, reward logic, disclosure timing, and response transparency. These dimensions were selected because they recur across vulnerability disclosure guidance and empirical studies of disclosure practice. Scope and reporting channel determine whether researchers can understand how to participate. Safe harbour determines whether participation appears legally safe. Reward logic determines whether bounty programmes look fair and predictable. Disclosure timing and response transparency determine whether the organisation accepts reciprocal obligations after a report is submitted. The approach follows the governance emphasis of the NIST Cybersecurity Framework 2.0, where cybersecurity outcomes are linked to organisational roles, policies, oversight, and risk management practices (NIST, 2024, Govern Function).

The analysis was interpretive rather than statistical. The aim was not to count every word or to produce a sector-wide ranking. Instead, the study examined how policy clauses work together to create a practical posture toward external researchers. For example, a policy may contain safe-harbour language but weaken it through vague scope or extensive unilateral discretion. Similarly, a bounty table may appear transparent while reward adjudication remains unclear. This form of close reading is useful because public disclosure policies often compress legal caution, technical instruction, and trust signalling into short clauses.

4. Discussion and Findings

4.1 Scope as a signal of organisational readiness

Scope is one of the clearest indicators of policy maturity. A well-written scope statement identifies eligible domains, products, applications, or services, while also explaining important exclusions. This matters because ambiguous scope transfers risk to the researcher. A researcher who cannot determine whether a target is authorised may either avoid testing or proceed under uncertainty. The problem is amplified in modern software environments where vulnerabilities may appear in dependencies, integrations, or web application components rather than in a single clearly bounded asset. Ruohonen's (2021) findings on Python package vulnerabilities support this concern because they show that software risk often emerges from ecosystem relationships rather than isolated code. The strongest policies treat scope as a practical guide rather than as a legal shield. They specify in-scope assets, distinguish first-party systems from third-party systems, and provide instructions for uncertain cases. This is important because vulnerability discovery ecosystems involve multiple actors and assets, not only one researcher and one vendor (Zhao et al., 2015). A narrow or unclear scope may protect the organisation legally, but it can reduce the number and quality of useful reports. In contrast, a clear scope supports coordinated vulnerability disclosure by reducing interpretive uncertainty before testing begins.

Scope also reveals organisational confidence. A company that identifies systems clearly appears to have coordinated internally among technical, legal, and operational teams. A company that uses general language such as "our services" may preserve flexibility but gives less guidance to researchers. Comparative work on coordinated vulnerability disclosure policies suggests that this variation is common and that scope is one of the areas where policy quality differs most visibly (Disso et al., 2023). Therefore, scope should be treated as a governance signal, not merely as an inventory list.

4.2 Reporting channels and procedural clarity

A disclosure policy becomes useful only when it explains how to report. Basic procedural details include the reporting channel, required technical information, preferred proof-of-concept format, affected asset, vulnerability type, reproduction steps, impact description, and contact details. These requirements improve triage quality and reduce unnecessary communication. They also show that the organisation has prepared itself to receive vulnerability information systematically. This aligns with NIST's recommendation that vulnerability disclosure programmes define intake and handling processes rather than rely on ad hoc reporting (NIST, 2023). Weak policies often provide only a generic email address or platform link without explaining what the report should contain. This can disadvantage less experienced researchers and increase triage burden for the company. Stock et al. (2016) showed that vulnerability notification can fail even when technical information is available, partly because communication and recipient readiness are difficult. A public policy can reduce this friction by telling researchers how to make their reports actionable.

Procedural clarity also has a trust-building function. When a policy explains acknowledgement, validation, escalation, and communication expectations, it signals that reports will not simply disappear into a generic inbox. In threat intelligence terms, vulnerability reports are valuable only when they become timely and actionable inputs into defensive decision-making (Tounsi & Rais, 2018). A clear reporting procedure therefore supports both operational efficiency and researcher confidence.



4.3 Safe harbour and conditional protection

Safe harbour is central because researchers need to know whether good-faith testing will expose them to legal or disciplinary action. Strong safe-harbour language typically states that the organisation will not initiate legal action against researchers who comply with the policy, avoid privacy intrusion, minimise data access, avoid service disruption, and report promptly. Zolkipli et al. (2022) describe safe harbour as a governance issue because it defines the conditions under which external cybersecurity work becomes legitimate rather than suspicious. However, many policies frame protection in highly conditional terms. Conditions are necessary because no organisation can authorise unrestricted experimentation on live systems. The problem arises when the conditions are vague, excessive, or one-sided. If a policy says that the company alone decides whether activity was authorised, or that any deviation removes protection, then the safe harbour may offer limited practical reassurance. O'Hare et al. (2024) show that organisational responses to coordinated vulnerability disclosure involve a continuing tension between trust and control. Safe harbour is where that tension becomes most visible. A mature policy should therefore distinguish harmful misconduct from minor, non-harmful overreach that can occur during legitimate testing. It should tell researchers to stop when they encounter sensitive data, report immediately, and avoid further access, but it should also avoid implying that accidental contact automatically turns good-faith research into hostile conduct. This balance supports responsible reporting because it makes legal protection credible without removing the company's ability to protect users and systems.

4.4 Reward structures and hidden discretion

Bug bounty policies are often the most visible expression of external security engagement because reward tables are concrete and easy to publicise. However, reward visibility does not necessarily mean reward transparency. Finifter et al. (2013) found that vulnerability reward programmes can produce useful reports, but programme effectiveness depends on design and administration. Maillart et al. (2017) similarly show that bug bounty outcomes depend on how programmes structure participation, not simply on the number of available researchers. The difficulty is that reward decisions often involve hidden discretion. A report may be rejected as a duplicate, downgraded because exploitability is considered low, excluded because the asset is out of scope, or denied because the issue was already known. These decisions may be legitimate, but policies often explain them only briefly. Munaiah and Meneely (2016) argue that the relationship between severity and bounty decisions can be inconsistent, which matters because researchers use rewards as signals of how organisations value their work. If the policy does not explain core adjudication principles, researchers may perceive the programme as unpredictable even when the company believes it is acting reasonably.

Bug bounty programmes also operate within market-like environments. Vulnerability markets contain information asymmetries because researchers know the discovery process while companies control validation and payment (Kannan et al., 2016). Allodi (2017) further shows that vulnerability trade and exploitation are shaped by economic factors beyond the technical severity of the flaw. Public bounty policies should therefore address not only maximum payouts but also eligibility, duplicates, severity assessment, report quality, and appeal or clarification mechanisms. Without these details, the reward table may create symbolic transparency while leaving the actual decision process opaque.

Platform governance adds another layer. Laszka et al. (2018) show that validation in bounty platforms can involve misaligned incentives, while Ruiz et al. (2021) describe the institutionalisation of bug bounty governance through platform-mediated rules and practices. This means a company using a platform should not rely entirely on the platform's existence as proof of fairness. Its own public policy should still explain how the organisation interprets eligibility, impact, communication, and reward decisions.

4.5 Disclosure timing, remediation, and reciprocity

Most disclosure policies ask researchers to report privately and delay public disclosure until remediation is complete or coordinated permission is granted. This request is often reasonable because premature public release can increase user risk. However, delayed disclosure is fair only when the organisation accepts reciprocal responsibilities. Formal disclosure models show that timing affects incentives and social welfare, so disclosure delay should not be treated as an unlimited company privilege (Arora et al., 2008 ; Kesan et al., 2019). In many policies, the researcher's obligations are more specific than the company's obligations. Researchers must avoid harm, keep findings confidential, submit detailed evidence, and wait for approval before disclosure. The company may only promise to review reports "as quickly as possible." This phrase preserves flexibility but provides little accountability. Walshe and Simpson (2020) demonstrate that vulnerability disclosure economics depends on incentives and response behaviour, which means that vague company commitments can weaken the disclosure relationship. A stronger policy would include acknowledgement windows, validation expectations, update intervals, and an approach to coordinated public disclosure where remediation is delayed. These do not have to be rigid deadlines in every case. Complex vulnerabilities may require investigation, customer coordination, and staged remediation. Still, a mature policy should give researchers enough information to interpret silence and understand whether progress is occurring. Organisational response behaviour is part of the disclosure process, not an internal matter completely separate from the policy.



4.6 Researcher identity, tone, and legitimacy

Policy tone matters because many security researchers operate outside formal organisational relationships. They must infer from public wording whether the company sees them as contributors, tolerated outsiders, or potential threats. Respectful language can support legitimacy by recognising good-faith research as a contribution to resilience. However, tone cannot substitute for substance. A policy can sound collaborative while preserving one-sided discretion, or it can be plain but operationally stronger. The institutionalisation of vulnerability discovery has changed the researcher-company relationship. Householder and Manion (2011) describe vulnerability discovery as increasingly industrialised, while Tchernin et al. (2020) frame bug bounty programmes as markets for responsible disclosure. These developments mean that researchers are not merely occasional volunteers; they may be skilled participants in a structured security economy. Policies that treat them only as risky outsiders may discourage experienced researchers who expect professional process and clear rules. At the same time, researcher legitimacy is not automatic. Organisations must protect customers, systems, privacy, and service availability. The challenge is to draft policies that distinguish good-faith research from harmful conduct without making legitimate research feel presumptively suspicious. A balanced tone should be respectful, precise, and firm. It should welcome useful reports, explain boundaries, and avoid excessive legalism where operational guidance would be more helpful.

4.7 Third-party assets and shared infrastructure

Technology companies increasingly depend on cloud services, third-party platforms, embedded components, open-source packages, and software supply chains. Public disclosure policies often respond by excluding third-party assets unless explicitly listed. This is understandable because companies cannot authorise testing of systems they do not control. Nevertheless, rigid exclusions can create reporting gaps because many real vulnerabilities appear at boundaries between services. Research on vulnerability ecosystems and software dependencies supports this concern. Zhao et al. (2015) describe web vulnerability discovery as an ecosystem involving multiple participants, while Ruohonen (2021) shows that software package vulnerabilities are embedded in wider application environments. If policies define responsibility too narrowly, researchers may be unsure how to report vulnerabilities involving authentication integrations, cloud storage, shared APIs, or inherited components. Mature policies should therefore include boundary guidance. They can state that researchers must not test third-party systems without authorisation, but they can also provide escalation instructions where a vulnerability affects users through a company-controlled service. This would preserve legal boundaries while acknowledging technical reality. A policy that treats third-party assets only as exclusions may be easier to draft, but it may fail to address the way modern technology systems actually operate.

4.8 Symbolic openness versus operational openness

A recurring pattern across public vulnerability disclosure and bug bounty policies is the difference between symbolic openness and operational openness. Symbolic openness appears in the existence of a security page, a reporting email, a bounty table, or a hall of fame. Operational openness appears in clear scope, credible safe harbour, usable reporting guidance, transparent triage, and reciprocal communication commitments. The two can overlap, but they are not the same. This distinction is important because research on vulnerability rewards and markets shows that participation depends on more than the availability of payment. Researchers consider fairness, expected response, legal risk, duplicate handling, and the likelihood that their work will be valued (Finifter et al., 2013; Tchernin et al., 2020). A company that advertises a high maximum bounty but provides vague rules may attract attention while still weakening trust. Conversely, a company with modest rewards but clear process may appear more credible to serious researchers. From a governance perspective, operational openness is the stronger indicator of maturity. It aligns with NIST's view that cybersecurity governance requires documented roles, processes, policies, and oversight (NIST, 2024, Govern Function). It also aligns with vulnerability disclosure guidance that emphasises preparation, coordination, and communication (NIST, 2023). A policy should therefore be judged by how well it structures the relationship between company and researcher, not by how visible or polished it appears.

5. Proposed Characteristics of a Mature Policy

The analysis supports a more demanding model of vulnerability disclosure policy maturity. First, a mature policy should define scope clearly. It should identify authorised assets, explain exclusions, and provide instructions for ambiguous cases. Scope should reduce uncertainty rather than merely protect the organisation from liability. This is especially important in ecosystem-based software environments where vulnerabilities may cross product, dependency, or third-party boundaries (Ruohonen, 2021; Zhao et al., 2015). Second, a mature policy should provide practical reporting instructions. Researchers should know what channel to use, what evidence to include, how to handle sensitive data, and whether secure communication is supported. This helps convert vulnerability information into actionable security intelligence, which is consistent with the broader value of timely and credible cyber threat information (Tounsi & Rais, 2018). It also reduces unnecessary triage friction for the organisation. Third, a mature policy should include credible safe-harbour language. It should promise not to pursue legal action against researchers who act in good faith and follow the policy, while clearly identifying harmful conduct that is not protected. The policy should avoid turning every minor uncertainty into a potential violation. Safe harbour is not only a legal clause; it is a trust mechanism that



determines whether researchers feel safe enough to report (O'Hare et al., 2024; Zolkipli et al., 2022). Fourth, a mature policy should explain response expectations. Acknowledgement windows, validation updates, remediation communication, and coordinated disclosure expectations make the relationship more reciprocal. Vendor response behaviour and patch timing are central to disclosure outcomes, as shown by studies of patch release and disclosure timing (Arora et al., 2010; Kesan et al., 2019). A policy that demands patience from researchers should also explain how the organisation will earn that patience. Fifth, a mature bug bounty policy should clarify reward logic. It should explain eligibility, severity evaluation, duplicate handling, report quality expectations, and reasons why otherwise valid findings may be excluded. This is necessary because bounty programmes operate within incentive systems that can easily become misaligned (Laszka et al., 2018; Munaiah & Meneely, 2016). Clear reward logic does not eliminate discretion, but it makes discretion more accountable.

Finally, a mature policy should align tone and substance. Collaborative language should be supported by operational commitments. A policy that praises researchers while preserving unlimited discretion may appear performative. A credible policy can still be firm, but it should respect researchers as participants in a shared security process. This reflects the broader movement toward professionalised and institutionalised vulnerability discovery (Householder & Manion, 2011; Ruiz et al., 2021).

6. Conclusion

This paper has analysed public vulnerability disclosure and bug bounty policies as cybersecurity governance documents. The analysis shows that these texts do more than provide contact information. They define permissible behaviour, allocate risk, frame researcher legitimacy, establish incentive structures, and reveal how organisations balance collaboration with control. Across the themes examined, a consistent pattern emerges: companies often invite external security research while retaining strong authority over scope, legal protection, reward eligibility, triage, remediation, and disclosure timing. This pattern does not make disclosure policies ineffective by definition. Organisations face legitimate responsibilities involving service availability, customer privacy, legal compliance, and third-party infrastructure. They cannot invite unrestricted testing. However, the documentary choices made in public policies matter because they influence whether researchers perceive reporting as safe, fair, and worthwhile. Prior research on vulnerability markets, reward programmes, disclosure timing, and vendor response behaviour shows that the success of disclosure depends on institutional design as much as technical discovery (Allodi, 2017; Arora et al., 2008; Finifter et al., 2013; Savage et al., 2015).

The central argument is therefore that policy maturity should be judged by documentary quality rather than visibility alone. A visible policy page, a branded bounty portal, or a general statement of appreciation cannot substitute for clear scope, credible safe harbour, transparent reporting instructions, reciprocal response expectations, and fair reward logic. Where these features are present, the policy becomes a practical instrument of cybersecurity governance. Where they are absent, the policy may remain largely symbolic even if the organisation appears open to external reporting. Future research could extend this study by comparing policies across sectors such as finance, healthcare, education, and public administration. It could also combine document analysis with interviews involving researchers, programme managers, legal counsel, and platform operators. Longitudinal studies would be valuable because vulnerability disclosure policies often change in response to incidents, regulatory pressure, or shifts in community expectations. Even without those extensions, this paper shows that public vulnerability disclosure and bug bounty documents are analytically rich texts that deserve closer attention in cybersecurity policy research.

References

- Algarni, A., & Malaiya, Y. K. (2014). Software vulnerability markets: Discoverers and buyers. *International Journal of Computer, Electrical, Automation, Control and Information Engineering*, 8(3), 449-456.
- Allodi, L. (2017). Economic factors of vulnerability trade and exploitation. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 1(1), 1-24.
- 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.
- Arora, A., Telang, R., & Xu, H. (2008). Optimal policy for software vulnerability disclosure. *Management Science*, 54(4), 642-656.
- Disso, J. P. P., Jones, K., & Bhandari, S. (2023). Coordinated vulnerability disclosure policies in practice: A comparative analysis. In *IEEE Conference on Communications and Network Security Proceedings* (pp. 1-6).
- Finifter, M., Akhawe, D., & Wagner, D. (2013). An empirical study of vulnerability rewards programs. In *USENIX Security Symposium Proceedings* (pp. 273-288).
- Householder, A. D., & Manion, A. (2011). The industrialization of vulnerability discovery in software security. *IEEE Security & Privacy*, 9(2), 7-10.



- Kannan, K., Telang, R., & Xu, H. (2016). Market for software vulnerabilities? Think again. *Management Science*, 62(6), 1709-1727.
- Kesan, J. P., Hayes, C. R., & Bashir, M. (2019). Value of software vulnerability disclosure timing. *Computer Law & Security Review*, 35(5), Article 105336.
- Laszka, A., Zhao, Y., & Grossklags, J. (2018). Banishing misaligned incentives for validating reports in bug bounty platforms. In *International Conference on Decision and Game Theory for Security Proceedings* (pp. 1-19).
- Maillart, T., Zhao, M., Grossklags, J., & Chuang, J. (2017). Given enough eyeballs, all bugs are shallow? Revisiting Eric Raymond with bug bounty programs. *Journal of Cybersecurity*, 3(2), 81-90.
- Munaiah, N., & Meneely, A. (2016). Vulnerability severity scoring and bounties: Why the disconnect matters. *IEEE Security & Privacy*, 14(5), 80-83.
- National Institute of Standards and Technology. (2023). Recommendations for vulnerability disclosure programs. National Institute of Standards and Technology.
- National Institute of Standards and Technology. (2024). Cybersecurity framework 2.0. National Institute of Standards and Technology.
- O'Hare, R., Whelan, C., & Sammon, D. (2024). Organisational responses to coordinated vulnerability disclosure: Process, trust and control. *Journal of Cyber Policy*, 9(1), 33-52.
- Ruiz, A., Winter, P., & Lindskog, S. (2021). HackerOne and the institutionalisation of bug bounty governance. In *Workshop on the Economics of Information Security Proceedings* (pp. 1-15).
- Ruohonen, J. (2021). An empirical analysis of vulnerabilities in Python packages for web applications. *Journal of Information Security and Applications*, 58, Article 102749.
- Savage, S., Coogan, P., & Lau, H. (2015). The economics of vulnerability disclosure. *ACM Computing Surveys*, 47(3), 1-38.
- Shahzad, M., Shafiq, M. Z., & Liu, A. (2012). A large scale exploratory analysis of software vulnerability life cycles. In *International Conference on Software Engineering Proceedings* (pp. 771-781).
- Stock, B., Pellegrino, G., Rossow, C., Johns, M., & Backes, M. (2016). Hey, you have a problem: On the feasibility of large-scale web vulnerability notification. In *USENIX Security Symposium Proceedings* (pp. 1015-1032).
- Tchernin, O., Chothia, T., & Novakovic, C. (2020). Bug bounty programs as markets for responsible disclosure. *Computers & Security*, 95, Article 101851.
- Tounsi, W., & Rais, H. (2018). A survey on technical threat intelligence in the age of sophisticated cyber attacks. *Computers & Security*, 72, 212-233.
- Walshe, T., & Simpson, A. (2020). An empirical study of the economics of vulnerability disclosure. *Information and Computer Security*, 28(1), 1-20.
- Zhao, M., Grossklags, J., & Liu, P. (2015). An empirical study of web vulnerability discovery ecosystems. In *CCS Workshop on Security Information Workers Proceedings* (pp. 13-24).
- Zolkipli, M. F., Jantan, A., & Jamil, N. (2022). Safe harbour in cybersecurity disclosure practice: A governance perspective. *International Journal of Advanced Computer Science and Applications*, 13(9), 412-420.