"People were not happy when admitting mistakes": Evolving Past Defensiveness in Running Coordinated Vulnerability Disclosure Programs

Sandra Rivera Pérez
Delft University of Technology
s.l.riveraperez@tudelft.nl

Amy Tjin Tjoen Jin
Delft University of Technology
a.m.tjintjoenjin@tudelft.nl

Michel van Eeten
Delft University of Technology
m.j.g.vaneeten@tudelft.nl

Carlos H. Gañán
Delft University of Technology
c.hernandezganan@tudelft.nl

Abstract

Coordinated Vulnerability Disclosure (CVD) is widely recommended as a security best practice and is increasingly mandated by regulation, yet adoption remains uneven. We report findings from 18 semi-structured interviews with security professionals from 15 organizations across multiple sectors and countries, including organizations with mature CVD programs, recent adopters, and organizations without formal CVD programs. We examine whether pre-adoption concerns identified in prior work persist, change, or materialize in practice, and how vulnerability handling and post-adoption challenges evolve over time.

Our findings show that commonly cited pre-adoption concerns—such as high report volume, poor-quality reports, and limited internal capacity—do occur, but usually in weaker and more manageable forms than anticipated. Several concerns emphasized in earlier studies, including lack of reporter expertise and difficulties maintaining reporter engagement, do not appear in our data. Instead, we identify post-adoption challenges that reflect program maturity, including handling differences between internal and external reports and responding to increasing regulatory pressure. By comparing organizations with and without CVD programs, we show that non-adoption is shaped by organizational capacity, leadership support, and regulatory awareness rather than a single barrier. Overall, our results suggest that CVD programs have become more normalized and less disruptive in practice, while remaining shaped by evolving organizational and regulatory demands.

1 Introduction

Coordinated Vulnerability Disclosure (CVD) has become a cornerstone of contemporary software security practice [1]. By providing a structured mechanism for receiving, triaging, and disclosing vulnerability reports from external parties, CVD programs are widely promoted as a way to improve product security, reduce adversarial disclosure, and foster constructive engagement with the security research community. As a result, CVD has transitioned from a voluntary "best practice" to an increasingly institutionalized requirement: governments and regulators now explicitly expect vendors to maintain public disclosure channels and internal vulnerability-handling processes [2–5].

Despite this normalization, adoption remains uneven. A 2025 report by the Internet of Things (IoT) Security Foundation found that only 40% of global IoT vendors have implemented a CVD program [6]. Particularly small and medium-sized organizations still lack formal CVD programs, even as regulatory pressure intensifies. This persistent adoption gap raises a basic but unresolved question: what challenges do organizations actually face when running CVD programs today, and how do these experiences differ from the concerns that shape adoption decisions in the first place?

Over the past decade, both academic and practitioner communities have developed a rich understanding of vulnerability disclosure. Prior work has documented operational challenges in running CVD programs, including concerns about report volume and quality, internal coordination costs, and tensions with vulnerability reporters [7–10]. Qualitative studies of disclosure from the reporter's perspective further highlight unresponsive communication, delayed remediation, and adversarial interactions, while organizational studies describe fears of legal exposure, reputational harm, and resource strain [7, 11–16].

Importantly, existing organizational studies—most notably work examining firms that had recently launched or were operating CVD programs—suggest that many of these challenges persist over time. However, the prevailing empirical picture is largely static: it implicitly assumes that the challenges identified during early adoption remain equally salient as programs mature, scale, and become embedded in organizational routines.

At the same time, the context in which CVD operates has changed substantially. Many vendors now have years—or decades—of experience with disclosure programs. Internal security teams, vulnerability-handling workflows, and third-party platforms have professionalized. Most critically, regula-

tory frameworks increasingly make vulnerability handling a legal obligation rather than a discretionary choice.

These changes also require distinguishing between two different forms of change examined in this paper. First, we consider how concerns evolve within organizations before and after adopting CVD programs. Second, we compare our present-day findings with challenges emphasized in earlier CVD literature, much of which reflects earlier stages of CVD adoption and ecosystem development. This distinction allows us to separate persistent operational challenges from those that appear to weaken, normalize, or evolve over time.

These shifts motivate our research question: *Do the pre- and post-adoption challenges emphasized in earlier CVD research still meaningfully constrain organizations today, or have they become normalized, mitigated, or replaced by new forms of friction?*

In this paper, we revisit Coordinated Vulnerability Disclosure from an organizational perspective, explicitly situating CVD within a landscape of program maturity, organizational learning, and regulatory pressure. We report findings from 18 semi-structured interviews with security professionals from 15 organizations across multiple sectors and countries, spanning a continuum from no formal CVD program to programs that were adopted less than a year ago to mature disclosure operations that have been running for up to 15 years.

Our findings reveal a consistent pattern: many pre-adoption fears do materialize in practice, but in substantially weaker and more manageable forms than suggested by earlier studies. Concerns such as high report volume, poor-quality submissions, and limited internal capacity are real but rarely threaten the viability of CVD programs once organizations gain experience. At the same time, challenges emphasized in prior work – such as lack of reporter expertise or difficulties sustaining engagement with researchers – are notably absent from our data.

Instead, we identify a set of post-adoption challenges that reflect program maturity rather than fragility, including handling discrepancies between internally and externally discovered vulnerabilities, navigating internal defensiveness around public disclosure, and adapting CVD practices to expanding regulatory requirements. These findings suggest that CVD programs have become more normalized and less disruptive than earlier research implies, while simultaneously becoming entangled with new organizational and legal constraints.

By contrasting organizations with and without CVD programs, we further show that non-adoption is rarely explained by a single barrier. Rather, adoption depends on the convergence of multiple enabling factors – technical capacity, organizational openness, leadership support, and regulatory awareness – while non-adopters tend to face deficits in one or more of these dimensions.

This paper makes three contributions: (1) an updated empirical account of CVD operations that captures how established challenges persist, attenuate, or disappear over time; (2) a comparative perspective that incorporates organizations without CVD programs to illuminate barriers to adoption; and (3) new insights into emerging challenges driven by regulatory pressure and organizational scale. Together, our results reposition CVD not as an inherently fragile practice, but as an increasingly routine organizational capability—one that most organizations can sustain, yet not without ongoing institutional effort.

2 Background and Related Work

Householder et al. define CVD programs as "the process of collecting vulnerability information, coordinating its distribution among stakeholders, and publicly disclosing vulnerabilities and mitigations" [1]. CVD programs provide a structured way for organizations to engage with external security researchers and manage vulnerabilities responsibly. Two common types of CVD programs are bug bounty programs, in which organizations offer monetary rewards for eligible vulnerability reports under specified conditions, and Vulnerability Disclosure Programs (VDPs), also known as responsible disclosure, which do not offer monetary rewards but may provide other forms of recognition, such as public acknowledgment in a hall of fame [8].

In this section, we review prior work on challenges in implementing and operating CVD programs and the countermeasures proposed from both the organizational (Section 2.1) and reporter (Section 2.2) perspectives. We also summarize the current legal frameworks and established best practices related to CVD programs (Section 2.3).

2.1 Challenges and Countermeasures from the Organization's Perspective

Prior work identifies recurring challenges for organizations when adopting and operating Coordinated Vulnerability Disclosure (CVD) programs [7, 8, 10]. Walshe and Simpson [8], based on a survey of 39 security professionals and eight follow-up interviews, provide a detailed account of both pre-launch and post-launch challenges.

Before launching a CVD program, organizations express concerns about receiving large volumes of low-quality or duplicate vulnerability reports, exposure to legal and reputational risks, and difficulties in communicating with external security researchers [7, 8]. Distrust of hackers and uncertainty about their motivations further intensify these concerns [17–19]. Internal constraints—such as limited security expertise, staffing shortages, and restricted budgets—also contribute to hesitation in adopting CVD programs [8, 10].

After adoption, many of these concerns continue to appear in practice. Organizations report recurring issues with poor-quality reports, submissions outside program scope, inflated severity assessments, and high report volume [7, 8, 20]. Prior work often attributes these problems to low-skilled or

poorly motivated reporters, especially in open programs that lack effective filtering mechanisms [21]. To address this, earlier studies recommend clearer scope definitions, private or invitation-only programs, and structured review processes to reduce noise and improve report quality [8, 21]. Recent work on AI vulnerability disclosure further highlights these scope-definition challenges. A large-scale study of AI vendors found substantial variation in how organizations define and handle AI-related vulnerabilities, with issues such as jailbreaking and hallucinations frequently treated as out of scope, while more traditional security concerns are more consistently accepted [22]. Some work also proposes automated tools for filtering and triaging reports [23, 24]. However, practitioners remain cautious about relying on automated report management systems, mainly due to concerns about their reliability in real-world settings [25].

Organizations also face challenges in their interactions with vulnerability reporters. Studies report communication breakdowns, disagreements over severity assessments, motivation-related conflicts, and undesirable reporter behavior [8, 9]. These challenges are influenced by differing incentive structures, including tensions between monetary and non-monetary motivations [13, 26]. Internally, organizations struggle with coordination between security and product teams, resistance to prioritizing remediation, and shortages of skilled personnel [7, 10]. Prior work identifies effective internal communication, managerial support, and a mature security culture as key countermeasures to these organizational challenges [7, 10, 27, 28].

2.2 Challenges and Countermeasures from the Reporter's Perspective

A complementary body of qualitative research examines the challenges faced by vulnerability reporters when engaging with organizations [7, 11–13]. These studies consistently show that reporters experience unresponsive or dismissive communication, delayed remediation, and, in some cases, legal threats. Prior work attributes these challenges in part to organizations' failure to clearly communicate disclosure expectations, legal constraints, and remediation timelines [29].

Research also documents differing views on disclosure practices between reporters and organizations. Many reporters—especially experienced ones—prefer full disclosure and choose not to collaborate with vendors, believing that public disclosure, sometimes accompanied by detailed exploits, applies pressure on organizations to fix vulnerabilities [13]. However, empirical evidence shows that vulnerabilities not reported to vendors are less likely to be fixed [13, 30].

Beyond interview-based research, several studies report first-hand experiences of researchers attempting to disclose vulnerabilities to specific organizations [14–16]. These works describe a wide range of organizational responses, from constructive engagement to silence or outright rejection. As a result, some reporters disengage from disclosure programs and become inactive over time [31]. van Hove et al. [16] show that many organizations remain difficult to contact about security issues, even when a dedicated security contact point exists. They also find that organizations with a formal VDPs respond more often and resolve more vulnerabilities. However, even among these organizations, roughly half of disclosed vulnerabilities remain unanswered or unresolved after 90 days, indicating that having a policy is necessary but not sufficient for effective vulnerability handling.

2.3 Legal Framework and Best Practices

Policy and regulation increasingly frame CVD programs as an expected organizational capability. In the European Union, the NIS2 Directive establishes a CVD framework in which Member States designate a Computer Security Incident Response Team (CSIRT) to coordinate CVD and support a European vulnerability database [2]. More recently, the EU Cyber Resilience Act makes vulnerability handling enforceable for vendors of digital products and requires manufacturers to implement a CVD program as part of vulnerability management obligations [5, 32].

Outside the EU, U.S. federal policy also treats disclosure handling as a formal requirement. CISA's Binding Operational Directive 20-01 requires U.S. federal civilian agencies to develop and publish a VDP and maintain supporting vulnerability handling procedures [4]. In the United Kingdom, government guidance similarly recommends that software and digital services provide an accessible vulnerability disclosure process backed by a clear policy and internal handling procedures [3].

Alongside regulation, best-practice guidance describes what a "well-functioning" CVD includes in practice. Standards such as ISO/IEC 29147 define recommended practices for vulnerability disclosure, while ISO/IEC 30111 specifies processes for internal handling of reported vulnerabilities (e.g., intake, triage, remediation, and coordination roles such as a PSIRT[1]) [33, 34]. At the European level, ENISA publishes guidance on CVD, including recommendations for harmonized approaches and national vulnerability coordination policies [35]. In the U.S., CISA provides practical guidance and templates for constructing a VDP, including expected elements such as a reporting channel, scope, response expectations, and safe reporting terms [4, 36].

Together, these regulations and standards [2–5, 33–36] establish a set of expectations: organizations are expected to provide a public reporting channel, define disclosure scope and timelines, handle vulnerability reports through structured internal processes, and coordinate remediation and disclosure. In general, these frameworks require vulnerabilities to be addressed in a timely manner, rather than within a fixed deadline such as the commonly cited 90-day window, which stems from long-standing coordinated disclosure norms, most

[1] *Product Security Incident Response Team*

prominently associated with Google's Project Zero disclosure policy [37].

3 Methodology

To examine the current state of CVD programs, we conduct semi-structured interviews with 18 participants from 15 organizations, including three organizations without CVD programs, between February 2025 and January 2026. This section presents our methodology: participant recruitment (Section 3.1), study design and interview protocol (Section 3.2), data analysis (Section 3.3), and ethical considerations (Section 3.4).

3.1 Recruitment and Participants

Participants were recruited from organizations across a diverse range of sectors (Table 2). Eligible participants were professionals involved in product security activities related to vulnerability disclosure, security governance, vulnerability discovery, disclosure handling, remediation processes, disclosure policy development, management of incoming reports, or coordination of vulnerability fixes within their organizations. Recruitment relied on two complementary strategies. First, we used existing professional networks and prior research contacts to identify relevant stakeholders and facilitate introductions to appropriate personnel within organizations. These individuals were then contacted with information about the study objectives and invited to participate. Second, we conducted a broader outreach campaign targeting organizations both with and without public CVD programs. Specifically, we contacted 25 organizations with publicly available CVD policies using the contact information listed in their disclosure programs. To include organizations without formal disclosure policies, we additionally contacted 48 companies through direct email and LinkedIn outreach. Organizations without public CVD programs were identified using data from the IoT Security Foundation report [6].

In total, 18 professionals from 15 organizations participate in the study, including 12 organizations with CVD programs and three without. Each participant and organization is assigned a unique identifier: a participant ID (PID) and an organization ID (OID). Participant demographics are presented in Table 1.

3.2 Study Design and Interview Protocol

We begin the study with an exploratory phase that includes two pilot interviews informed by the research question, prior literature on CVD programs, and an initial interview protocol. These interviews clarify how organizations manage vulnerability disclosure in practice, the steps involved, and the challenges encountered, and they inform revisions to the protocol. After these first two interviews, we make minor adjustments to the wording and sequence of some questions while retaining the overall structure and content. The final interview protocol is presented in Appendix A.

The interview protocol consists of two comparable sections: one for organizations with formal CVD programs and one for organizations without such programs. This structure captures differences in vulnerability management practices and organizational perspectives while maintaining consistency across interviews. During data collection, we make minor adjustments to question wording and order while preserving the overall structure to follow emerging themes when necessary.

After obtaining informed consent (see Section 3.4), each interview begins with general questions about the participant's role, responsibilities, and organizational context. The protocol then follows a structured order aligned with prior work, particularly the study by Walshe and Simpson [8], allowing analysis of whether organizational perspectives on CVD programs change over time. We first address pre-adoption considerations and organizational views on CVD. Ideally, these pre-adoption considerations would be discussed with participants involved from the beginning of the program to capture initial views and early-stage processes. However, this is not always feasible. As shown in Table 1, seven participants are involved in setting up the CVD program. For participants not involved from the start, pre-adoption questions are adapted to focus on the organization's current motivations and concerns involving CVD programs, shifting the emphasis from its early-stage to its present state.

We then examine vulnerability discovery and remediation practices, including reporting, assessment, and resolution processes. Next, we explore interactions with vulnerability reporters, and we conclude with questions on post-adoption challenges, countermeasures, and recommendations. When these topics are not raised spontaneously, we probe for challenges identified in prior CVD studies and for current organizational practices. This approach enables comparison between present-day CVD operations and challenges reported in earlier research.

All interviews are conducted remotely using video conferencing tools. Most interviews involve one participant and the researcher; one interview includes two participants and the researcher (P16–O14 and P17–O14). Seventeen interviews are conducted in English, and one in Spanish (P18–O15). Interviews last an average of 44 minutes, with durations ranging from 23 to 59 minutes.

3.3 Data Analysis

We analyze the interviews using thematic analysis and continue data collection until reaching theoretical saturation, defined as the point at which no new meaningful themes emerge [38, 39]. All interviews are recorded and transcribed. After anonymization and translation of the Spanish interview, the primary researcher analyzes the transcripts using open

Table 1: Participant Demographics

PID	OID	Country	Role	Experience within the organization	With CVD?	Interviewee present when CVD started
P1	O1	NL	Product security lead	6–10 years	✓	×
P2	O1	NL	Product security and policy	>20 years	✓	✓
P3	O2	US	Security engineer	2–5 years	✓	×
P4	O2	US	Incident manager	11–15 years	✓	×
P5	O3	NL	Security analyst	11–15 years	×	–
P6	O4	NL	Threat management	2–5 years	✓	×
P7	O5	NL	Security officer	2–5 years	✓	×
P8	O6	DE	Security engineer	6–10 years	✓	✓
P9	O7	NL	Security officer	6–10 years	✓	✓
P10	O8	NL	Product security lead	2–5 years	✓	×
P11	O9	NL	Security officer	≤1 year	×	–
P12	O10	NL	Security officer	6–10 years	✓	✓
P13	O11	NL	Product security lead	2–5 years	✓	✓
P14	O12	DE	Product security lead	>20 years	✓	✓
P15	O13	US	CEO	2–5 years	×	–
P16	O14	NL	Security incident responder	6–10 years	✓	×
P17	O14	NL	Security officer	>20 years	✓	✓
P18	O15	ES	Offensive security specialist	2–5 years	✓	×

Note. Participants report whether they were present during the establishment of the CVD program.

coding, annotating emerging themes and developing an initial codebook. A single coder conducts this stage of analysis, which is appropriate for this form of qualitative research [40].

The codebook is then discussed with two other researchers with varying backgrounds, and the codes are refined to better fit the data. This process of validation and refinement continues as the analysis progresses. Once theoretical saturation is reached, the final codebook is applied to all interview transcripts to ensure consistent coding across the dataset.

Although full saturation is not reached for organizations without CVD programs due to the smaller sample size (three participants), we include these findings to provide complementary perspectives and to offer insight into current vulnerability handling practices and barriers to CVD adoption, which remain underexplored in prior work.

3.4 Ethics and Data Protection

The study is approved by the human research ethics committee of the researchers' institution. Before each interview, participants receive written informed consent and are also informed orally by the primary researcher about the purpose of the study, the voluntary nature of participation, the absence of compensation, and the data collection process. All transcripts and quoted materials are anonymized. Audio recordings, transcripts, and interview data are stored on a secure network at the researchers' institution.

4 Results

In this section, we address our research question: *Do the pre- and post-adoption challenges emphasized in earlier CVD research still meaningfully constrain organizations today, or have they become normalized, mitigated, or replaced by new forms of friction?* To answer this question, we present findings on pre-adoption considerations, including motivations and concerns (Section 4.1); how vulnerability management operates in practice (Section 4.2); post-adoption challenges and the countermeasures used to address them (subsubsection 4.3.1); and participant recommendations for adopting and operating CVD programs (Section 4.4).

Because prior work identifies organizational characteristics as key drivers of security practices [41, 42], including CVD programs [7, 8], we first characterize the 15 organizations included in our study. Table 2 provides an overview of these organizations, including country, size (based on Gartner's definition [43]), sector, the presence or absence of CVD, and the estimated duration of the CVD program as reported by the participants.

4.1 Pre-adoption Considerations

Here, we present the motivations (subsubsection 4.1.1) and concerns (subsubsection 4.1.2) reported by organizations at different stages of CVD adoption when deciding whether to implement or not implement a CVD program.

Table 2: Overview of Organizations

OID	Country	Size	Age	Sector	With CVD?	CVD duration
O1	NL	Large	>50 years	Manufacturing	✓	11–15 years
O2	US	Large	31–50 years	Manufacturing	✓	>20 years
O3	NL	Small	11–30 years	Technology	×	–
O4	NL	Large	31–50 years	Finance	✓	6–10 years
O5	NL	Medium	>50 years	Technology	✓	2–5 years
O6	DE	Large	>50 years	Manufacturing	✓	6–10 years
O7	NL	Medium	≤10 years	Technology	✓	≤1 year
O8	NL	Large	11–30 years	Manufacturing	✓	2–5 years
O9	NL	Small	11–30 years	Education	×	–
O10	NL	Large	>50 years	Education	✓	6–10 years
O11	NL	Large	11–30 years	Retail	✓	2–5 years
O12	DE	Large	>50 years	Manufacturing	✓	11–15 years
O13	US	Small	≤10 years	Technology	×	–
O14	NL	Large	>50 years	Technology	✓	11–15 years
O15	ES	Large	11–30 years	Finance	✓	2–5 years

Organization size is according to Gartner's definition [43].
The duration of CVD programs is based on information provided by the participants.

4.1.1 Motivations

Based on the interviews, we identify five motivations that led organizations to implement CVD programs: leveraging external security expertise, maintaining control over disclosure, building trust and protecting reputation, meeting legal requirements, and responding to market incentives.

External expertise. The first motivation, mentioned by 11 participants, is to improve security by involving external researchers, a motivation also identified by Alomar et al. [7] in the context of bug bounty adoption. One participant describes CVD as a way "to make our products even more secure by allowing external people to look at our products and assess them," because it is "helpful if we have external security resources looking at our systems and reporting vulnerabilities so we can improve" (P1-O1). Another participant describes this as having "a second set of eyes to identify the external exposure and the vulnerabilities and weaknesses" (P6-O4). One participant links this motivation to the growth of the research community, noting that "the number of security researchers has dramatically increased in the past 7–8 years" (P10-O8). Another participant highlights the value of reach, stating that "the biggest advantage is access to global talent, even well-funded companies cannot hire top specialists in every niche" (P18-O15).

Disclosure control. The second motivation, mentioned by nine participants, is the desire to be involved in vulnerability disclosure and to serve as the first point of contact for reporters. Two participants explain that "CVD programs open the gates to people to find us" (P13-O11), so that reporters "report it first to us and then we can work together on bringing this to the news, but allowing us first to fix it and inform our customers" (P1-O1). CVD is also seen as a way to provide "a clear channel of communication" and make it "as easy as possible for people that notice anything wrong in our products to communicate that to us" (P9-O7), giving researchers "a space to come to you rather than to first go public" (P12-O10). A participant also stresses the importance of shaping disclosure outcomes, noting that "we also want to be involved in the narrative" (P1-O1).

Trust and reputation. The third motivation, mentioned by eight participants, is maintaining customer trust and protecting reputation. One participant describes CVD programs as "doing the right thing" and part of the "need to do good for the customers" (P4-O2). Another participant notes that CVD helps "customers trust the reputation the organization has built" (P1-O1), while another points out that "if it came out that our product is not secure, it would crash our reputation" (P5-O3). One participant also links CVD to responsibility toward users, emphasizing the need to "make sure that we have a safe network and infrastructure for ourselves, but also for our customers" (P16-O14).

Legal requirements. The fourth motivation relates to legal aspects and is mentioned by four participants. One participant explains that CVD programs help by "lowering your exposure to lawsuits and that kind of thing" (P4-O2), while another points out that CVD offers "a more formal way of handling vulnerabilities, especially when reports involve legal liabilities" (P13-O11).

Market incentives. The fifth motivation relates to market incentives for implementing CVD programs and is mentioned by two participants. One participant notes that "security and fast remediation is something that you can sell ... something that you can market" (P5-O3). Another participant describes adopting CVD because it became expected in the industry, stating that their company "needed to have a CVD as basically other companies already had it" (P14-O12).

Participants also describe two motivations for not imple-

menting CVD programs: competing organizational priorities and a perception that CVD is an unnecessary expense because the organization faces low perceived security risk.

Organizational priorities. One participant from an organization without a CVD program explains that adoption was discussed internally but they "decided to not go with it yet," noting that "the adoption is delayed due to competing priorities, restructuring, and a lot of mergers and acquisitions. We have to put our house in order before we do that [implementing CVD]" (P11–O9).

Unnecessary. Another participant from an organization without a CVD program, where no prior discussions had taken place, describes CVD as unnecessary due to the perceived low risk of vulnerabilities. They explain that their systems handle limited or low-value data, so "there isn't much data to be useful, and even if accessed, it's an annoyance, but it's not going to cause harm to our customers" (P15–O13). The same participant also notes that responsibility for sensitive data and security processes lies with a third-party platform, stating that "they're maintaining all the data and systems ... we don't really keep anything, they have to deal with that topic [security, CVD]" (P15–O13).

4.1.2 Concerns

When analyzing the pre-adoption concerns mentioned by participants, six concerns emerge: volume of reports, poor-quality reports, negative press, distrust of reporters, internal capacity, and process liability.

Volume of reports. The first concern relates to the expected volume of reports. Four participants, all from organizations with CVD programs, describe fears that launching a CVD program would mean "opening the floodgates to a lot of reports," requiring effort "to kind of go through the report [and] evaluate if it's something that is super urgent or if it's something that can wait" (P13-O11).

Poor-quality reports. The second concern, mentioned by four participants from organizations with CVD programs, relates to the quality of reports. Participants anticipate receiving reports for issues they are already aware of or consider low priority, such as "a vulnerability in a certain version of a certain library that is low severity, and not something I would spend time fixing" (P8-O6).

Negative press. The third concern, reported by three participants from organizations with CVD programs, relates to negative press. One participant describes fears that vulnerabilities disclosed through CVD could "become known too quickly" and "lead to negative press, even when we follow due process and fix the issue before disclosure" (P1-O1). The same participant also explains the concern that "after publishing an advisory, a day later there's a huge article in the news saying the organization has lots of problems, which then requires explanation to leadership because they see the negative press" (P1-O1). Another participant reports resistance to public disclosure, noting that some preferred to "not have public statements about vulnerabilities" due to concerns about the company's public image and the belief that "we don't have vulnerabilities, we have robust products" (P14-O12).

Distrust of reporters. The fourth concern, reported by two participants from organizations with CVD programs, is distrust toward reporters and relates to fears that CVD programs could encourage hacking of organizational systems and retaliation. One participant reports pushback from regulatory and legal teams, noting that "the last thing that we want is to encourage people to start hacking our systems," and emphasizing that "we do not want to motivate people to start hacking our products" (P2–O1). Another participant describes concerns that CVD could be seen as "inviting bug bounty hunters or unintended traffic on our production systems" (P13–O11).

Internal capacity. The fifth concern relates to internal capacity to implement a CVD program and is raised by two participants from organizations without CVD programs. Participants explain that limited financial and human resources make adoption difficult, particularly for smaller organizations. One participant notes that "a small company can't afford to go hire somebody ... it doesn't financially make any sense" (P15–O13). They add that if CVD "became regulated by law, we would have to go hire a firm and pay some money, which would be just an exercise of money being spent" (P15–O13). Another participant emphasizes the lack of manpower, explaining that "we don't have the manpower at the moment to create an official process and do anything really good with it, even though we do pay attention to security in our own way" (P5–O3).

Process liability. The sixth concern, reported by one participant from an organization without a CVD program, relates to the risk of having a program that is not properly followed. The participant explains that adopting CVD requires "a process that you need to follow, and the one thing worse than not having a program at all is having a program that [you are] not following at all" (P5–O3). They further emphasize that "if on paper you have a process, but if in practice you don't do anything with it, the program just creates additional risk" (P5–O3).

Table 3 compares pre-adoption concerns identified in prior work—often referred to as fears—with those reported in our study. The qualitative "strength" indicators (limited, medium, strong) are based on how frequently themes were discussed during the interviews and how many participants mentioned a particular concern. The dot indicators reflect the relative prominence of issues and should not be interpreted as quantitative measurements. Prior studies were mapped onto the same scale based on the findings reported by the authors. Since we did not have access to the original interview data from these studies, this mapping relied on how strongly particular themes were emphasized in their reported findings and discussion. This allows for approximate comparisons across studies while acknowledging methodological and interpretive differences

between them.

Overall, Table 3 shows that five pre-adoption concerns identified in prior studies persist in our findings, but their perceived impact is lower. Concerns that were previously described as strong—such as report volume, report quality, and distrust of reporters—are still recognized but are generally viewed as manageable and unlikely to prevent adoption. Another strong concern in prior work, internal capacity, is perceived as manageable by larger organizations but remains a more significant concern for smaller organizations. Concerns about negative press appear more prominent in our findings than in prior work, which helps explain why a key motivation for adopting CVD programs is to be involved in vulnerability disclosure and shape the public narrative. Other concerns, including lack of reporter experience and communication with reporters, are not identified in our study as pre-adoption concerns. Finally, we identify process liability as a new concern not emphasized in prior work, raised by an organization without a CVD program.

4.2 Vulnerability Discovery and Remediation in Practice

Based on our findings, vulnerability discovery and remediation in practice can be described as a process consisting of four main phases, reflecting how participants experience and manage vulnerabilities. These phases are: intake and triage—covering both internal findings and external reports (subsubsection 4.2.1); assessment and prioritization (subsubsection 4.2.2); remediation (subsubsection 4.2.3); and public disclosure (subsubsection 4.2.4).

Two of these phases—intake and triage, and remediation—have been discussed to some extent in prior work by [7] in the context of bug bounty programs. More broadly, the four phases identified in our study align with elements of the disclosure process model described by [1], who note that they "adapt a version of the ISO/IEC 30111 process with more phases to better describe what we have seen at the CERT/CC." They are also consistent with prior studies of the vulnerability lifecycle [44–46].

Table 4 provides an overview of the CVD characteristics of the organizations in our study, which we reference throughout this section.

4.2.1 Intake and Triage

The intake and triage process differs between organizations with and without CVD programs. For organizations without a CVD program, the lack of a clear point of contact limits control over how vulnerability reports are received. As one participant explains: "without a formal channel, we don't have a lot of influence about how we get our report" (P5-O3).

Among organizations with CVD programs, "reports come mainly from two sources, internal or external" (P4-O2). A participant adds that "most of them [internal reports] come from engineers that can be part of customer support or from the business units themselves," such as "developers that find out something that looks like it might have security impact on the product" (P4-O2).

How external reports are received depends on how the CVD program is implemented (see Table 4). For organizations running CVD programs in-house, a common approach is to provide a dedicated email address. One participant describes having "a centralized e-mail inbox, with the public e-mail on the website," explaining that "that inbox is monitored during business hours and also after business hours" (P9-O7).

Once a report is received, organizations typically begin with an acknowledgment. Participants emphasize that this step is important to show that reports are handled by a real person. One participant explains, "upon receiving a vulnerability report, we basically make an acknowledgment first. So thank you for reporting," noting that this is important because when someone writes to an organization's e-mail, "you never know if you will ever get a response ... often this is just a black hole and it disappears" (P14-O12). The participant stresses that the response is manual: "we write these emails manually... there might be typos... there might be questions following. So that people see there's really someone sitting there, not just a script or an AI" (P14-O12). After acknowledgment, teams perform a quick triage: "is the claim reasonable? Is enough material provided?" (P14-O12). If not, they "ask back... can you provide more material or we don't understand what you did" (P14-O12).

For organizations using third-party platforms (e.g., Bugcrowd[2], HackerOne[3]), intake takes a different form. Internal security teams receive notifications when a report is submitted. As one participant explains, "every time a new report comes in, we get an e-mail notification and we also get the status of the report" (P10-O8). In these cases, "the triage... has usually happened externally" (P10-O8). However, internal teams still review the reports themselves. One participant explains that "the first assessment of the platform is correct I would say maybe 9 out of 10 times, but sometimes external reviewers don't know our company as well as we do," which can lead to severity being misjudged (P8-O6). As a result, "we also check all the technical descriptions... and I also assess the actual impact by ourselves" (P8-O6).

Reports received through channels outside the CVD program are uncommon. Participants note that "certain researchers try to reach out on an individual basis on social media," but in those cases organizations "redirect them to the formal process" (P6-O4). Some indirect reports still occur, as "we still get frequent reports via [other] company e-mails ... through social media, not so much" (P8-O6), or when "they reach out to our customer care service and then customer care sends it back to us" (P10-O8). Participants emphasize that

[2] https://www.bugcrowd.com/
[3] https://www.hackerone.com/

Table 3: Pre-adoption Concerns: Prior Studies vs. Our Findings

Pre-adoption concern	Prior studies	Our study	Interpretation from our study
Volume of reports [7, 8, 20]	●●●	●●	Recognized as a concern, but not expected to prevent adoption because the workload is expected to decrease over time.
Poor-quality reports [7, 8, 20]	●●●	●●	Recognized as a concern, but outweighed by expected useful reports.
Distrust of reporters [7, 8, 17]	●●●	●	Raised as a concern, especially by legal or non-technical stakeholders.
Internal capacity [7, 8, 20]	●●●	●	Recognized as a concern, particularly for smaller organizations.
Lack of reporter experience [8, 20]	●●	–	Not identified as a pre-adoption concern in our study.
Communication with reporters [8]	●●	–	Not identified as a pre-adoption concern in our study.
Negative press [8]	●	●●	Recognized as a concern, but not a barrier for adoption.
Process liability	–	●	Recognized as a concern by an organization without CVD programs.

Strong issue: ●●● Medium: ●● Limited: ● Not identified: –

Table 4: CVD Characteristics and Vulnerability Handling

OID	CVD type	Incentive	# internal vulnerabilities	# external vulnerabilities	Resolution timeline
O1	In-house	Hall of fame	30–40 per year	10 per year	Ideally < 60 days
O2	Hybrid	Ack. Sometimes monetary	100–200 per year	40 per year	Variable
O3	No CVD	–	10 per year	3 per year	Variable
O4	Third-party	Sometimes monetary	More than reported	1–2 per week	Hours to months
O5	In-house	No reward	500–1000 per year	1–3 per year	3–30 days
O6	Third-party	Hall of fame	80–90%	∼10%	36 hours to months
O7	In-house	No reward	∼1–2 per year	Mostly external (90%)	1–7 days
O8	In-house	Hall of Fame	1 per week	1 per month	2 weeks to 3 months
O9	No CVD	–	40–50 per year	6 per year	–
O10	In-house	Hall of fame	∼1 internal to ∼10 external	∼1 internal to ∼10 external	Variable
O11	In-house	Other	500–1000 per year	30–40 per year (10% valid)	Variable
O12	Hybrid	Ack. Sometimes monetary	300 per year	50 per year	Variable
O13	No CVD	–	No vulnerability so far	No vulnerability so far	–
O14	In-house	Hall of Fame. Other	4000–5000 per year	50 per year	Variable
O15	Third-party	Monetary	More than reported	2000 (100 valid) per year	3–7 days

Hybrid: CVD programs are primarily managed in-house but use third-party platforms for selected products.
Ack: Reporters are acknowledged in published advisories.
Sometimes monetary: Monetary rewards are offered only in specific cases or for selected products.
Other: Non-monetary rewards are offered (e.g., T-shirts).
 # vulnerabilities and timeline: Values are interviewee estimates quoted directly from the interviews.

"it is very rare to actually see a vulnerability being posted on the Internet that we don't know about" (P10-O8). A few cases involve conference disclosures—"there's cases where they [reporters] go to conferences and we find out like that" (P10-O8)—but overall, as one participant summarizes, "no full public disclosures have occurred, and only in rare cases researchers hinted at findings on social media without details" (P18-O15).

Overall, intake and triage follow a similar sequence: receiving the report, acknowledging the reporter, conducting an initial assessment of validity and completeness, and assigning the report to the next stage of the vulnerability handling process, namely the internal assessment.

4.2.2 Assessment and Prioritization

Assessment and prioritization follow intake and triage and involve a deeper analysis of the reported vulnerability. While triage checks basic validity, this phase focuses on understanding the impact, context, and urgency.

Participants explain that assessment depends on how and where the affected component is used. One notes that "it really depends on which component is impacted by the vulnerability," and that even "a CVSS[4] score of 10 could be 0 in our case" depending on implementation and existing controls (P1-O1). Reports are also verified for reproducibility, as "sometimes it can be a misconfiguration ... or we cannot reproduce the issue" (P2-O1).

From what participants report, each organization uses its own way to assess criticality and prioritize vulnerabilities, even when common metrics such as CVSS are used. One participants describes that "as one of the first steps, we assess severity using CVSS v3.1. But, in most cases, we deviate from it because there's not enough granularity in the CVSS to really explain what's going on" (P4-O2). Organizations adjust severity based on their internal context and risk models, using predefined levels and timelines. Some organizations combine metrics, saying that "we've prioritized them using two different metrics ... one is the CVSS, the other one is the EPSS[5]" (P13-O11). Others rely on internal models, noting that "we have a prioritization model based upon the risk and that gives you the time to solve it" (P17-O14).

Differences between external and internal assessments are common. Participants note that researchers may report vulnerabilities that "are technically valid but whose real-life exploitability is nearly impossible," making them "not a high priority" internally (P2-O1). External reporters may also "misjudge some reports because they simply don't know our company, requiring teams to reassess the actual impact" (P8-O6).

Based on this assessment, organizations assign severity, identify the affected asset and owner, and create internal tickets for remediation. As one participant summarizes, after validation, they "assign severity, create an internal remediation ticket and forward the issue to the business unit" (P18-O15).

4.2.3 Remediation

Once a vulnerability report is assessed and assigned to the responsible developers or business units, the remediation process begins. These teams plan and implement the fix, and participants report significant variation in how this work is done across organizations and teams. As one participant explains, "some teams run completely on their own, and after receiving a vulnerability, two weeks later they come back and say, yeah, we will fix this, this is our timeline" (P14-O12). In other cases, remediation involves "closer cooperation [with security teams]... with regular meetings step by step throughout the process" to evaluate progress and check whether "the timeline is still valid or not" (P14-O12).

Remediation often starts by determining whether an immediate fix is possible or whether interim mitigations are needed. Participants note that "the ideal mitigation is ... a software update, but that it's not always easy, as fixes may require an architectural change," which can take more time (P1-O1). In such cases, teams may look for temporary measures to mitigate the risk, for example by changing settings or applying operational controls. Others describe blocking exposure at the network level, explaining that when "this thing does not have a fix within the software itself, we can use features to just block it" (P7-O5). Some organizations isolate affected systems entirely, stating that "if the system is not business-critical, we take it down, and if it is business-critical, we put it behind the firewall" (P12-O10). These interim measures are used until a permanent fix can be deployed.

Remediation timelines vary widely, as shown in Table 4, and depend on the severity of the vulnerability and the complexity of developing a permanent fix. Participants emphasize that "it's really difficult to make statements on how quickly that is, since remediation can take anything from weeks to months" (P2-O1). Critical issues are handled fastest, with fixes reported "within four to six hours" or "within the same day ideally" (P9-O7; P13-O11). Less severe issues follow longer timelines, commonly "two weeks to three months" (P10-O8). Several organizations define explicit timelines by severity, such as "72 hours for critical issues," with lower severities handled within "seven days, 14 days or 30 days" (P7-O5). As one participant summarizes, "each rating has a timeframe to act upon it" (P6-O4).

Overall, remediation is described as a flexible process that depends on severity, complexity, and organizational structure. Organizations combine permanent fixes with temporary mitigations to reduce risk while remediation is ongoing. Once fixes or mitigations are in place, organizations move to the next phase of the process: public disclosure.

4.2.4 Public Disclosure

After remediation, organizations communicate vulnerability information and fixes to customers through public disclosure. Participants report using multiple channels, depending on the customer type, product, and legal obligations. Most organizations publish a public security advisory; as one participant explains, "we're publishing our advisories on our website," and "if it comes into CVD programs ... there will of course be an advisory" (P4-O2; P2-O1). These advisories are "completely public" (P1-O1).

Many organizations complement public advisories with direct customer notifications. Participants describe customer portals where users can log in and see "these are my devices, that's where I'm impacted ... these are the fixes," making the process "more transparent" (P1-O1). Others use subscription-based updates, where customers "subscribe to newsletters ... and they get an update like that," often filtered by the products they use (P1-O1). Some organizations rely on more direct communication, such as "e-mail or phone calls... the old fashioned way," or send information through "release notes,

[4] *Common Vulnerability Scoring System*
[5] *Exploit Prediction Scoring System*

newsletters, or monthly reports" (P5-O3; P9-O7; P11-O9).

Communication practices also vary by customer size and relationship. Larger customers may have "dedicated communication channels with the security teams," while smaller customers are reached through "mailing lists" (P8-O6). In business-to-business settings, disclosure may follow "contractual agreements" and go "directly to the responsible person" (P10-O8). In some cases, legal requirements apply, and organizations must issue "field safety notices ... where we really directly inform the customers" (P2-O1).

Overall, public disclosure is handled through a mix of public advisories and targeted communication, and is shaped by customer needs, regulatory requirements, and organizational practices.

4.3 Post-adoption Challenges and Countermeasures

In this section, we discuss the post-adoption challenges reported by participants in our study (subsubsection 4.3.1) and the countermeasures organizations use to address these challenges (subsubsection 4.3.2).

4.3.1 Post-adoption Challenges

Here, we discuss twelve distinct post-adoption challenges mentioned by participants in our study: validity of reports, handling variability, poor-quality reports, internal coordination, volume of reports, internal capacity, staffing and cost, interaction with reporters, leadership understanding and support, defensiveness from internal teams, public exposure, and regulatory pressure. We examine whether the pre-adoption concerns (see subsubsection 4.1.2) actually materialize in practice, and we compare our findings with challenges identified in prior work—referred to as issues in previous studies—to assess which challenges persist, disappear, or emerge, and to what extent they affect CVD programs.

Validity of reports. Disagreements about whether a report describes a real vulnerability and how severe it is are a common post-adoption challenge. Seven participants explain that "there has been cases where we would, for instance, not agree on the criticality," but emphasize that this is "normal" and not viewed as a major problem (P2–O1). Participants explain that they have not had "very bad experience, except those cases where we really would not agree on the criticality or the severity of the vulnerability" (P2–O1).

These disagreements often arise because "the researcher usually doesn't have a lot of knowledge about the software we use," which can lead to claims that cannot be reproduced or exploited in practice (P5–O3). One participant gives an example where researchers claimed "you have a SQL injection vulnerability," but internal checks showed that "we haven't been able to really make use of the vulnerability," because "we also have an extra check ... that handles all of those injections" (P5–O3). As a result, researchers may have "a limited view on what our software or what our part of the software does, making it hard for them to understand why we say no, this isn't a vulnerability" (P5–O3). Participants describe these as "very vivid discussions," but explain that they are usually resolved through communication, such as "having a call with the researchers and explain to them why, in our view, something isn't a vulnerability," or by relying on third parties "to clear up that noise and pass to us whatever is really validated," helping avoid prolonged debate (P5–O3; P6–O4).

Handling variability. Handling variability between internally and externally reported vulnerabilities is identified as a post-adoption challenge by five participants. Internally discovered vulnerabilities are generally described as easier and smoother to handle, as organizations have earlier visibility and clearer ownership. As one participant explains, "in our internal security process, we know already much earlier when we have a vulnerability ... we can plan and we can organize according to that. Once identified, issues are placed into our normal life cycle management process and released through planned updates to customers" (P4–O2). Another participant adds that internally, "everything is very well streamlined ... we know which systems are processing these problems, we know the people who are working on this because we have ticketing systems, mail trails, and logs" (P11–O9). A third highlights that when a vulnerability is found internally, "you are already talking to an expert," and "the time to respond is way quicker and way more reliable because you already have the right people involved" (P10–O8).

External reports, in contrast, often introduce additional pressure and extra steps. Participants explain that "the moment something comes in through CVD, the clock starts ticking," because public disclosure becomes a concern, which "drives some steps in the CVD that we don't have on the internal" (P2–O1). Others note that external findings "become a must to fix, even when they disrupt normal development work" (P6–O4). Overall, participants describe this challenge as stemming from the contrast between streamlined internal processes and externally reported vulnerabilities, which add time pressure and coordination effort.

Poor-quality reports. Four participants mention report quality as a challenge and note that it "varies a lot among reporters" (P4–O2). Several point out that low-effort submissions occur at times, explaining that "sometimes it happens ... that the reports are completely garbage" (P4–O2). Others emphasize that "some reports are poorly written or contain unprofessional language, which creates extra work" (P18–O15). At the same time, participants consistently stress the value of well-prepared reports. One explains that "the more thorough the research and the report is, the better" and points out that "the most important thing is to understand which version they are running, what is the configuration of the gear they are using ... and then also clear reproduction steps, because as soon

as we have reproduction steps, then we can put that to the test in our lab as well" (P4–O2). Others describe high-quality submissions where reporters "made a great report on how they found it and how we could reproduce it," sometimes "even including guidance on how to fix it" (P9–O7). Participants also observe changes over time, noting that "there were quite good reports when we started," while "the spam we see [is] more recently now that time has passed" (P13–O11). Overall, participants conclude that although low-quality reports create some extra work, high-quality submissions provide clear value and make report quality a manageable issue rather than a barrier.

Internal coordination. Internal coordination emerges as a post-adoption challenge, mentioned by three participants, especially in large and complex organizations. Participants explain that scale and diversity make vulnerability handling difficult, as " we are a large organization, many business units, many different products," and when someone reports "I found a problem in product X," teams first need to understand "which business unit is responsible for product X" (P2–O1). One participant mentions that, unlike smaller organizations where "it could even be one person who does it all," larger ones involve "a lot of people in different departments," which makes it hard to consistently "find the responsible group and report it there" (P17–O14). Coordination is further complicated by the fact that teams have "different backgrounds, different academic studies, yet are expected to do everything in a similar fashion, in a similar way," which "can be a difficult situation sometimes" (P17–O14). Even when formal processes exist, participants note that "finding the right communication line and finding the correct colleagues" remains a challenge that "doesn't reduce the time, and increases the effort quite a lot" (P10–O8). Overall, participants highlight that handling reports internally is often less about the technical fix and more about navigating organizational size, multiple communication paths, and limited visibility across teams, which together make coordination slow and effort-intensive.

Volume of reports. Across organizations, participants consistently report that vulnerability volume is dominated by internal reports rather than external ones, as shown in Table 4. The majority of organizations (12 out of 15) indicate that most vulnerabilities are identified internally, with estimates such as "75 to 80%" or "80 to 90%" of findings coming from internal sources, compared to only a small number of external reports (P4-O2; P7-O5). Only two organizations (O7 and O10) report a higher volume of external than internal findings, while one organization (O13) reports neither internal nor external vulnerabilities. One participant explains that "it's above 90%, but most of the time it's in the software we use. It's not our own software, but it's in the IT platforms and the application server, web service, all those third party software we use," and adds that "we recently started a security training with internal employees, so that might change in the near future because they also get the tools and knowledge to discover and do something about security" (P9–O7).

These results address a common pre-adoption concern that CVD programs would lead to an unmanageable influx of external reports. Participants acknowledge that "there was a lot of submissions" when programs were first launched, but note that "that huge demand goes back to normal, like one or two per week" (P6-O4). As one participant summarizes, "you can expect at the beginning a lot of tickets going to you, but afterwards you should be able to reduce the number and make it within the manageable capacity" (P6-O4). Participants further explain why internal findings dominate, noting that "internally, we find more issues because we have more context and access. Internal testing is closer to grey-box or white-box testing, while external researchers work in black-box conditions. When external researchers find something critical, it often has significant merit because they are working without internal knowledge" (P18-O15).

Internal capacity, staffing, and cost. Internal capacity, staffing, and cost emerge as post-adoption challenges, particularly in relation to organizational size and structure. Participants note that sustaining a CVD program requires ongoing operational effort, including "ensur[ing] there is sufficient staffing" and having "sufficient people looking at the inbox," which is described as "also a challenge" (P2–O1). One participant from an organization without a CVD program highlights that these concerns are especially relevant for smaller and medium-sized organizations, where limited capacity can lead to basic failures, such as reports "ending up in spam and never being seen, or being deleted without proper assessment" (P11–O9).

Participants emphasize that capacity challenges go beyond headcount and include maintaining routine processes. Even small operational tasks require attention, as "you have a PGP key, but you need to keep that up to date," and forgetting to do so has happened "one or two cases over the years," even if considered "a minor thing" (P2–O1).

From a cost perspective, participants consistently report that CVD programs are not overly expensive. Several organizations rely on recognition instead of payments, noting that "we do recognition, but we don't do a reward or financial reward," and therefore "I don't think it's that costly to implement the process" (P1–O1). While organizations still need "a security person in your organization to be able to assess it," costs are generally manageable (P1–O1). Some participants observe that "smaller companies often assume CVD takes a lot of time and a lot of money. But in reality is not that expensive, since you just make a website, and you just structure it [CVD] well" (P2–O1). One participant even describes CVD as "a cheap way to have security testing" (P2–O1).

Decisions around third-party platforms are also shaped by cost and capacity. While some organizations prefer to stay in-house "to be in full control," others report that it was "actually cheaper to engage an external platform than to pay our internal team doing that work" (P8–O6). At the same time, participants

caution that "bug bounty programs… for smaller companies are relatively costly," highlighting that cost–benefit trade-offs depend on organizational size and maturity (P8–O6).

Overall, participants describe internal capacity, staffing, and cost as manageable post-adoption challenges, with constraints most visible in smaller organizations.

Reporter behavior. Participants describe interactions with reporters as mostly positive and cooperative. Communication "generally goes pretty well, with good alignment and understanding with security researchers" (P1–O1). Many participants stress that most reporters are trying to help improve security, and overall interactions are described as "very, very good," where "90 to 95% of the cases are positive and negative cases are rare" (P2–O1; P8–O6).

When problems occur, participants mainly point to two sources: disagreements about validity (discussed previously) and pressure from reporters (mentioned by three participants). Pressure most often relates to rewards or disclosure timelines. One participant describes a case where a reporter approached the organization by saying "I found this and I'm going to put it in public if you don't pay me," which "is something that the company doesn't like" (P11–O9). Another participant reports pressure around strict timelines, noting that some researchers are "really super adamant of giving us exactly 90 days," even though "sometimes it's just not realistic for us to be able to fix something in 90 days" (P4–O2). When reporters refuse to accept delays, they may publish "no matter what," which forces organizations to disclose as well, "even if we don't have fixes" (P4–O2).

Overall, participants describe interactions with reporters as largely positive, with challenges limited to occasional pressure around rewards or disclosure timelines rather than systematic conflict.

Leadership understanding and support. This post-adoption challenge is mentioned by three participants. One explains that CVD programs often need "a little bit of a push because there are people on the board that don't really understand how these things are working" (P11–O9). When leadership lacks this understanding, vulnerability disclosure can easily be seen "as a cost center only," rather than as something that adds value (P4–O2). In contrast, participants stress the importance of support from higher management, noting that "the fact that we do have people higher up that are aware of CVD and that are willing to invest in this kind of program is really helping us" (P4–O2). At the same time, even when security is broadly supported, CVD does not always receive explicit priority, as "even if you have them nailed down on police to support security, it's not necessarily that the CVD program gets a push," since leadership must constantly decide whether to "spend it on security or to make the product cheaper to compete on the market" (P14–O12). Overall, participants describe leadership support as uneven, with limited understanding at higher levels sometimes slowing down or constraining CVD efforts.

Defensiveness from internal teams. Three participants identify defensiveness from internal teams as a post-adoption challenge. One explains that vulnerability reports are often received with skepticism, noting that "usually they [developers] are very sceptical ... because someone is coming out of the blue in order to tell you that you have a problem with your software and this is not well received by developers," even though "it is very well received by security professionals like me because I am always looking to see if there is a problem, a real problem there," adding that "for the developers it is a little bit of a different story" (P11-O9). Another participant points to resistance from legal teams, stating that "legal is also very hesitant on publishing information on how we do certain things and to promise timelines and things like that" (P10-O8). A third highlights a broader issue, explaining that "people were not happy when admitting mistakes, admitting failure" (P14-O12). Overall, participants describe defensiveness as a recurring but limited challenge, mainly affecting interactions between security teams, developers, and legal stakeholders.

Public exposure. Three participants identify public exposure as a post-adoption challenge, closely tied to disclosure timing and researcher behavior. Participants explain that organizations are cautious about disclosure because they "don't want to be held hostage in quotes by the researchers … when they have to remediate an issue with their product" (P3–O2). A key concern is disclosing vulnerabilities before a fix is ready, as organizations "try to avoid as much as possible disclosing something when there is no fix available, because doing so would expose us and our customers" (P4–O2). Early disclosure without remediation "puts us and our customers in a very bad position," since it means "we're signaling to the world exposure to something for which there is no fix" (P4–O2). Maintaining communication is therefore seen as critical, as one participant stresses the need "to have good communication with the researchers, because if you don't, things like public disclosure will happen, and it will be even worse for your reputation" (P5-O3). Overall, participants describe public exposure as a sensitive but manageable challenge that depends largely on effective communication and coordination around disclosure timing.

Regulatory pressure. Regulatory pressure is identified as a post-adoption challenge by three participants, particularly in relation to newer legal frameworks that make vulnerability disclosure mandatory rather than voluntary. One participant explains that disclosure "has been voluntary, but now … in some domains it's a legal obligation," pointing to regulations such as the Cyber Resilience Act in Europe and the PSTI Act in the UK (P2–O1). While large organizations are generally seen as capable of absorbing these requirements, participants express concern about smaller and medium-sized organizations, noting: "I really wonder to what extent small and medium-sized organizations really will be able to put this [regulations] in place" (P2–O1).

Participants also highlight that regulatory pressure shapes how much information organizations are willing to disclose. One explains that to manage legal risk, organizations "publish the least amount as possible and keep up to date with the legislations," and that new requirements may force them to "go to legal to ensure disclosures both protect us and still follow the regulations" (P10–O8). In regulated sectors, pressure is even stronger, as "if you have a finance company, even if it is small, they are taking care of this kind of vulnerabilities because of the strict framework" (P11–O9). Overall, participants describe regulatory pressure as an increasing constraint that formalizes CVD practices but also raises concerns about legal risk, disclosure scope, and the ability of organizations to comply.

Our findings show that most pre-adoption concerns do materialize in practice, but usually in a weaker form (see Table 3). Concerns such as report volume, poor-quality reports, and internal capacity are often described as strong barriers in earlier studies, but in our study they mostly emerge as medium or limited challenges after adoption. Participants report an initial increase in submissions and occasional low-quality reports, but these challenges tend to decrease over time and become manageable through basic triage and experience. Other pre-adoption concerns identified in earlier studies, including lack of reporter experience and difficulties communicating with reporters, do not materialize in our findings. Overall, while pre-adoption concerns are not unfounded, they rarely turn into long-term or unmanageable problems once a CVD program is in place.

Table 5 compares our findings on post-adoption challenges with those reported in prior work. As in subsubsection 4.1.2, the qualitative "strength" indicators (limited, medium, strong) reflect how frequently issues were discussed and how many participants mentioned them. Prior studies were mapped onto the same scale based on the findings reported by the authors, enabling approximate comparison across studies.

Overall, we observe both continuity and change in post-adoption challenges. Core challenges such as report validity, internal coordination, and internal capacity continue to affect organizations, confirming that these issues persist after adoption. However, participants generally describe their impact as lower than suggested in earlier studies, with issues like poor-quality reports, report volume, and reporter behavior viewed as manageable rather than critical. At the same time, our study identifies challenges that receive little to no attention in prior work, particularly handling differences between internal and external reports and the increasing regulatory pressure. In contrast, challenges emphasized in earlier studies—such as maintaining long-term engagement with reporters—are not identified by our participants. Taken together, these findings suggest that while many post-adoption challenges persist, their impact is reduced through organizational learning, and new challenges emerge as CVD programs mature and operate in more regulated environments.

4.3.2 Countermeasures

Participants describe three main countermeasures to address post-adoption challenges: updates to the CVD program, internal organizational changes, and the use of external platforms. Not all organizations implemented countermeasures, and several participants explicitly report making no changes.

Updates to the CVD program. Mentioned by four participants and also identified as the most common countermeasure by [8], these updates are mainly small adjustments rather than major redesigns. One participant explains that "what we have done over the years is indeed improve the process," including decisions about "what steps to follow and who to involve" (P2–O1). Another adds that they "adjusted the remediation timelines from time to time, like how long it should take to fix critical or medium vulnerabilities," even though "roughly, the process hasn't changed too much" (P4–O2). Some organizations also adapted how they interact with reporters, noting that "we've learned kind of to cut the conversation short if we notice that it's not going in the right direction" (P13–O11).

Internal organizational changes. This countermeasure is mentioned by three participants and is also identified in prior work [8]. These changes focus on improving internal coordination and clarifying responsibility. One participant explains that after launching the program, "we had to make a list of all internet-facing websites, IP addresses, and who's hosting it, and what's running on it," because "that was not clear in the beginning, which made routing reports difficult" (P9–O7). Another participant describes changes in decision-making structures, noting that "we created a whole process flow … and decisions who will be informed when about what," which helped clarify escalation paths and responsibilities across teams (P2–O1).

Use of external platforms. This countermeasure is mentioned by one participant and is also discussed in prior work [8]. External platforms are used mainly to reduce internal workload through triage. A participant explains that before using an external platform, "the first assessment of our vulnerability [was] within our internal team, but this is not really the most efficient way, especially since platforms such as HackerOne or Bugcrowd, have their own triage team as well." As a result, the organization decided "to outsource this job, so that our colleagues would have more time left to do other work," because "the first assessment you can relatively easily outsource," while "everything behind that is not so easily outsourceable" (P8–O6).

No countermeasures. Four participants report not implementing any countermeasures. Two explain that "our general disclosure process … hasn't really changed that much" (P8–O6), or that "while the program has been useful to us, we have other priorities to focus on, at least for the coming year" (P13–O11).

Overall, countermeasures tend to be incremental and pragmatic rather than radical. Organizations mainly adjust in-

Table 5: Post-adoption Challenges: Prior Studies vs. Our Findings

Post-adoption challenge	Prior studies	Our study	Interpretation from our study
Poor-quality reports [7, 8, 20]	●●●	●●	Low-quality or incomplete reports occur but are considered manageable.
Validity of reports [7, 8, 20]	●●●	●●●	Disagreements about scope or severity arise but are typically resolved through assessment and discussion.
Volume of reports [8, 20]	●●●	●	Participants report an initial spike of external reports after adoption that becomes manageable after a few months. Most of the vulnerabilities are discovered internally.
Internal coordination [7, 8]	●●	●●	Identifying responsible teams and coordinating remediation is complex, especially in large organizations.
Internal capacity, staffing, and cost [8, 20]	●●	●●	Staffing and cost constraints are generally manageable for larger organizations but pose greater challenges for smaller ones.
Reporter behavior [8]	●●	●	Interactions are mostly positive; pressure around rewards or timelines occurs in isolated cases.
Leadership understanding and support [7]	●	●	Limited leadership understanding persists but does not block program operation.
Maintaining reporter engagement [20]	●	–	Participants do not report difficulties maintaining engagement; respectful treatment encourages reporters to submit more reports.
Defensiveness from internal teams [8]	●	●	Defensive reactions from developers and legal teams occur mainly early in program adoption and decrease over time.
Public exposure	–	●	Disclosure remains sensitive, particularly when timelines are driven by external reporters.
Handling variability	–	●●	External reports are less predictable and introduce more pressure than routine internal findings.
Regulatory pressure	–	●●	Regulatory requirements increasingly shape CVD operations and disclosure practices.

Strong issue: ●●● Medium: ●● Limited: ● Not identified: –

ternal coordination, refine existing processes, or selectively outsource triage, while some introduce no changes at all. Importantly, three of the four countermeasures identified by [8] remain relevant in our study, while the use of additional tooling is not mentioned by participants. The fact that four participants report no countermeasures suggests that, although CVD programs do face challenges, these challenges are often not severe enough to require countermeasures.

4.4 Recommendations from Participants

Nine participants explicitly recommend implementing a CVD program, emphasizing that its benefits outweigh the challenges when it is set up properly and supported internally. Participant recommendations cluster around three themes: establishing a clear process and expectations, ensuring internal readiness, and starting small.

Clear process and expectations. Mentioned by nine participants, establishing a clear process and expectations is the most frequently emphasized recommendation. Participants stress that organizations should "build up the processes and discuss it first with the people who will be involved so they can give advice" (P11–O9). One participant mentions the importance of "set[ting] up clear guidelines, like how quickly you're gonna respond to something" (P3–O2). In addition, participants emphasize the need for a well-defined scope, advising organizations to "first make sure that they have a clear scope, including scoping the external assets to be included in these kind of programs, and clearly classifying what is considered in scope" (P6–O4).

Internal readiness. This recommendation is mentioned by seven participants. They emphasize that organizations should have basic security capacity and visibility before opening external reporting channels. One participant recommends addressing internal vulnerabilities first: "If I have my house in order like this internally, then I would think to go also externally" (P11–O9). Having asset oversight is also critical, as organizations should "have a good oversight of your assets

and of who does what before you do so" (P10–O8).

Start small. Mentioned by five participants, especially in relation to smaller organizations, participants recommend beginning with simple disclosure mechanisms and expanding over time. One recommends "start[ing] with a very basic process for vulnerability disclosure, a security.txt document that is easy to find" (P13–O11). For smaller organizations, participants note that "they might consider to do it in-house before investing in more complex programs" (P8–O6).

Overall, participants recommend CVD programs as beneficial but caution that success depends on clear processes, internal readiness, and a gradual, manageable implementation.

5 Discussion

Our findings suggest that how organizations experience CVD in practice today differs in important ways from how it is described in prior work [8–10, 17–20]. Rather than indicating a failure of earlier research, this difference reflects a shift in context: the most recent qualitative study with organizational security experts dates back to 2022 [8], before regulatory pressure and organizational experience had fully reshaped CVD practice. Regulatory developments—most notably the EU Cyber Resilience Act [5]—have moved CVD from an optional best practice toward a normalized organizational capability. Including organizations without CVD programs allows us to capture this transition directly, showing not only why CVD feels less disruptive in practice, but also why adoption remains uneven despite normalization.

The motivations described by participants suggest that CVD is primarily framed internally as a risk-management practice rather than as a visibility or signaling mechanism. Organizations emphasize access to external expertise, control over disclosure timing, and the protection of trust and reputation as ways to reduce uncertainty around vulnerability handling, rather than to attract attention or recognition. Legal aspects and market incentives matter, but they appear to reinforce adoption rather than initiate it. In contrast, organizations without CVD programs are not rejecting disclosure as a principle; instead, adoption is postponed when perceived security risk, available capacity, and organizational priorities do not align. This suggests that non-adoption reflects situational readiness rather than resistance, helping explain why CVD remains unevenly adopted even as it becomes increasingly normalized.

A central contribution of this study is showing that many pre-adoption concerns persist but no longer function as meaningful constraints. While earlier studies describe strong pre-adoption concerns, we find that many of these concerns do not block program adoption. Notably, two concerns highlighted in prior work—lack of reporter expertise and difficulties communicating with reporters—are absent and external researchers are instead viewed as a valued source of technical insight. This shift suggests that organizational learning and repeated interaction have reduced early-stage uncertainty around working with external reporters.

At the same time, our findings surface a previously unidentified pre-adoption concern: process liability. For some organizations, the perceived risk lies not in running a CVD program, but in committing to a process they cannot reliably execute. This reframes non-adoption not as resistance, but as a judgment about operational readiness, highlighting why some organizations prefer no program over a symbolic one.

Post-adoption challenges discussed by participants largely overlap with those reported in prior work, but—similar to pre-adoption concerns—they are generally less severe and less disruptive than earlier studies suggest [7, 8, 20]. Core issues—report validity, internal coordination, and capacity—remain, yet are treated as routine operational work rather than threats to program viability.

In contrast to prior work, challenges related to report volume, poor-quality reports, and reporter behavior are described as limited or moderate. High report volume is mainly experienced in the first months after adoption and tends to stabilize over time. Poor-quality reports increase workload but are viewed as a normal outcome of opening disclosure to a broad and diverse group of reporters. Interactions with reporters are mostly constructive, and when tensions arise they tend to reflect misaligned incentives rather than breakdowns in the disclosure process. These tensions are closely linked to reporter motivations: while some report vulnerabilities out of personal interest or curiosity [13, 18], others are primarily driven by financial rewards [26]. This perspective contrasts with prior, reporter-centered studies, which emphasize negative disclosure experiences and sustained conflict [7, 11–13].

At the same time, our study identifies post-adoption challenges that are largely absent from earlier work and reflect the realities of mature CVD programs. One key challenge is the risk of public exposure, especially when disclosure timelines are driven by external reporters and fixes are not yet available, creating pressure that is organizational rather than technical. Participants also highlight the asymmetry between internal and external vulnerability reports: internally found issues follow predictable processes with early visibility and clear ownership, while externally reported vulnerabilities often arrive unexpectedly, trigger disclosure expectations, and require additional coordination under time pressure. A third emerging challenge is regulatory pressure, as new laws increasingly mandate disclosure practices and remediation timelines, reducing flexibility in how organizations manage vulnerabilities. Together, these challenges suggest that friction in CVD has shifted—from uncertainty about running a program to constraints created by visibility, external accountability, and legal obligations—indicating a move from early-stage operational concerns to institutional and regulatory ones.

In contrast, one challenge highlighted in prior studies—maintaining reporter engagement—does not appear in our data. Previous work shows that "around 80% of reports

come from 20% of reporters" [47], suggesting that once reporters engage with a program, they often continue to participate.

The limited use of countermeasures—and the fact that some organizations adopt none at all—suggests that CVD has become sufficiently routinized to operate without continuous intervention. Incremental adjustments replace structural change, reinforcing the interpretation of CVD as a stable organizational capability.

Taken together, these findings answer our research question by showing that earlier CVD challenges have not disappeared, but have largely become normalized, reduced in importance, or replaced by new forms of friction. CVD today appears less fragile and less disruptive than prior work suggests. Instead, its impact is increasingly shaped by organizational scale, coordination demands, and regulatory requirements. The central challenge is no longer whether CVD programs can be sustained, but how they are integrated into broader governance, compliance, and development practices over time.

Based on our findings, we offer recommendations for both organizations and regulators. For organizations and security teams, we suggest not overestimating the risks of adopting CVD, as common concerns such as report volume, report quality, and internal capacity are generally manageable in practice. In line with participants' advice, organizations should set up a clear, organization-wide CVD process with a defined scope, roles, timelines, and escalation paths, and ensure basic internal readiness before opening disclosure channels. Starting with a simple setup and improving it over time helps build experience, while keeping in mind that external reports often create more time pressure than internal findings. For regulators, our results indicate that CVD requirements are feasible when expectations are clear and realistic. Regulatory frameworks should allow flexibility in timelines and implementation, especially for smaller organizations, and focus on transparency, responsiveness, and process maturity rather than strict procedures.

6 Limitations and Future Work

The main limitation of this study is participant selection. We were not able to recruit many organizations without CVD programs, which limits insight into non-adoption perspectives. In addition, the sample likely reflects selection bias, as larger organizations with more mature security practices may have been more willing to participate. The sample is also weighted toward organizations operating in the Netherlands, which may limit the generalizability of the findings to other regulatory or organizational contexts. Another limitation concerns the comparative analysis with prior work. The qualitative "strength" indicators used throughout the paper are interpretive and based on how frequently themes were discussed during interviews and how many participants mentioned a particular issue, rather than on quantitative measurements. Furthermore, prior studies were mapped onto the same scale based only on their reported findings and discussion, since we did not have access to their original interview data. As a result, these comparisons should be interpreted as approximate and treated with caution given the methodological and interpretive differences between studies. Future work should focus on smaller organizations and those without CVD programs to better understand barriers to adoption and how regulatory pressure and resource constraints shape CVD decisions across different organizational contexts.

7 Conclusion

This paper contributes an updated, organization-centered view of CVD programs based on 18 semi-structured interviews with security professionals from 15 organizations, including both organizations with established CVD programs and organizations without. By including organizations without CVD programs, our study brings perspectives that have been largely absent from prior qualitative work and helps explain both adoption and non-adoption decisions. We show that many pre-adoption concerns reported in earlier studies—such as report volume, report quality, and internal capacity—do occur in practice, but typically in milder and more manageable forms. Several concerns emphasized in prior work, including lack of reporter expertise and difficulties maintaining reporter engagement, are not observed in our data. Post-adoption challenges such as report validity, internal coordination, and staffing persist, but rarely threaten program operation, while new challenges related to regulatory pressure and differences between internal and external reports are becoming more relevant. Overall, our findings suggest that CVD programs have matured, becoming more stable, predictable, and normalized in practice, and that anticipated risks are less disruptive than commonly assumed.

References

[1] A. D. Householder, G. Wassermann, A. Manion, and C. King, "The cert guide to coordinated vulnerability disclosure," Carnegie-Mellon University, Pittsburgh, PA, United States, Tech. Rep., 2017.

[2] European Union, "NIS2 Directive: securing network and information systems," 2022. [Online]. Available: https://digital-strategy.ec.europa.eu/en/policies/nis2-directive

[3] United Kingdom, "The Product Security and Telecommunications Infrastructure (Security Requirements for Relevant Connectable Products) Regulations 2023," 2023. [Online]. Available: https://www.legislation.gov.uk/ukdsi/2023/9780348249767

[4] CISA, "BOD 20-01: Develop and Publish a Vulnerability Disclosure Policy," 2020. [Online]. Available: https://www.cisa.gov/sites/default/files/bod-20-01.pdf

[5] European Union, "Regulation (EU) 2024/2847 (Cyber Resilience Act) - Annex I, Part II: Vulnerability Handling Requirements," Official Journal of the European Union, 2024. [Online]. Available: http://data.europa.eu/eli/reg/2024/2847/oj

[6] IoT Security Foundation, "The State of Vulnerability Disclosure in Global Consumer IoT in 2025," Tech. Rep., 2025. [Online]. Available: https://iotsecurityfoundation.org/wp-content/uploads/2026/01/The-State-of-Vulnerability-Disclosure-Usage-in-Global-Consumer-IoT-in-2025V8.pdf

[7] N. Alomar, P. Wijesekera, E. Qiu, and S. Egelman, "You've Got Your Nice List of Bugs, Now What?" Vulnerability Discovery and Management Processes in the Wild," in *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. USENIX Association, 2020, pp. 319–339. [Online]. Available: https://www.usenix.org/system/files/soups2020-alomar.pdf

[8] T. Walshe and A. Simpson, "Coordinated Vulnerability Disclosure programme effectiveness: Issues and recommendations," *Computers & Security*, vol. 123, p. 102936, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167404822003285#bib0049

[9] T. Walshe, "Supporting data-driven software development life-cycles with bug bounty programmes," Ph.D. dissertation, University of Oxford, 2023. [Online]. Available: https://ora.ox.ac.uk/objects/uuid:4a828bbb-8ff4-4cac-9e09-5699b30c6d52

[10] A. Y. Ding, G. L. De Jesus, and M. Janssen, "Ethical hacking for boosting IoT vulnerability management," in *Proceedings of the Eighth International Conference on Telecommunications and Remote Sensing*. New York, NY, USA: ACM, 2019, pp. 49–55.

[11] A. Happe and J. Cito, "Understanding Hackers' Work: An Empirical Study of Offensive Security Practitioners," in *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. New York, NY, USA: ACM, 2023, pp. 1669–1680.

[12] J. E. Noordegraaf and M. Weulen Kranenbarg, "Why do young people start and continue with ethical hacking? A qualitative study on individual and social aspects in the lives of ethical hackers," *Criminology & Public Policy*, vol. 22, no. 4, pp. 803–824, 2023.

[13] M. Hafiz and M. Fang, "Game of detections: how are security vulnerabilities discovered in the wild?" *Empirical Software Engineering*, vol. 21, no. 5, pp. 1920–1959, 2016.

[14] G. C. M. Moura and J. Heidemann, "Vulnerability Disclosure Considered Stressful," *ACM SIGCOMM Computer Communication Review*, vol. 53, no. 2, pp. 2–10, 2023.

[15] T.-H. Chen, C. Tagliaro, M. Lindorfer, K. Borgolte, and J. Van Der Ham-De Vos, "Are You Sure You Want To Do Coordinated Vulnerability Disclosure?" in *2024 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2024, pp. 307–314.

[16] K. van Hove, J. v. d. H.-d. Vos, and R. van Rijswijk-Deij, "Your Vulnerability Disclosure Is Important To Us: An Analysis of Coordinated Vulnerability Disclosure Responses Using a Real Security Issue," 2023.

[17] L. M. Tanczer, "50 shades of hacking: How IT and cybersecurity industry actors perceive good, bad, and former hackers," *Contemporary Security Policy*, vol. 41, no. 1, pp. 108–128, 2020.

[18] D. Votipka, R. Stevens, E. Redmiles, J. Hu, and M. Mazurek, "Hackers vs. Testers: A Comparison of Software Vulnerability Discovery Processes," in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 374–391.

[19] L. Follis and A. Fish, "State hacking at the edge of code, capitalism and culture," *Information, Communication & Society*, vol. 25, no. 2, pp. 242–257, 2022.

[20] M. Al-Banna, B. Benatallah, D. Schlagwein, E. Bertino, and M. C. Barukh, "Friendly Hackers to the Rescue: How Organizations Perceive Crowdsourced Vulnerability Discovery." *PACIS*, vol. 230, 2018.

[21] M. Al-Banna, B. Benatallah, and M. C. Barukh, "Software Security Professionals: Expertise Indicators," in *2016 IEEE 2nd International Conference on Collaboration and Internet Computing (CIC)*. IEEE, 2016, pp. 139–148.

[22] Y. Piao, J. Li, and D. W. Woods, "" Abuse Risks are Often Inherent to Product Features": Exploring AI Vendors' Bug Bounty and Responsible Disclosure Policies," *arXiv preprint arXiv:2509.06136*, 2025.

[23] Y. Fan, X. Xia, D. Lo, and A. E. Hassan, "Chaff from the wheat: Characterizing and determining valid bug reports," *IEEE transactions on software engineering*, vol. 46, no. 5, pp. 495–525, 2018.

[24] X. Xia, D. Lo, Y. Ding, J. M. Al-Kofahi, T. N. Nguyen, and X. Wang, "Improving automated bug triaging with specialized topic model," *IEEE Transactions on Software Engineering*, vol. 43, no. 3, pp. 272–297, 2016.

[25] W. Zou, D. Lo, Z. Chen, X. Xia, Y. Feng, and B. Xu, "How practitioners perceive automated bug report management techniques," *IEEE Transactions on Software Engineering*, vol. 46, no. 8, pp. 836–862, 2018.

[26] S. S. Malladi and H. C. Subramanian, "Bug Bounty Programs for Cybersecurity: Practices, Issues, and Recommendations," *IEEE Software*, vol. 37, no. 1, pp. 31–39, 2020.

[27] J. Smith, C. Theisen, and T. Barik, "A Case Study of Software Security Red Teams at Microsoft," in *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 2020, pp. 1–10.

[28] M. Hilton, N. Nelson, T. Tunnell, D. Marinov, and D. Dig, "Trade-offs in continuous integration: assurance, security, and flexibility," in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*. New York, NY, USA: ACM, 8 2017, pp. 197–207.

[29] T. Walshe and A. Simpson, "Towards a Greater Understanding of Coordinated Vulnerability Disclosure Policy Documents," *Digital Threats: Research and Practice*, vol. 4, no. 2, pp. 1–36, 2023.

[30] S. Rivera Pérez, M. van Eeten, and C. H. Gañán, "Patchy Performance? Uncovering the Vulnerability Management Practices of IoT-Centric Vendors," in *2024 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, 2024, p. 153. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/SP54263.2024.00154

[31] T. Walshe and A. Simpson, "A longitudinal study of hacker behaviour," in *Proceedings of the 37th ACM/SIGAPP Symposium on Applied Computing*, 2022, pp. 1465–1474.

[32] J. Ruohonen and P. Timmers, "Vulnerability Coordination Under the Cyber Resilience Act," 2025. [Online]. Available: https://arxiv.org/abs/2412.06261

[33] "ISO/IEC 29147:2018," 2018. [Online]. Available: https://www.iso.org/standard/72311.html

[34] "ISO/IEC 30111:2019," 2019. [Online]. Available: https://www.iso.org/standard/69725.html

[35] ENISA, "Vulnerability Disclosure," 2026. [Online]. Available: https://www.enisa.europa.eu/topics/vulnerability-disclosure

[36] CISA, "Vulnerability Disclosure Policy Template." [Online]. Available: https://www.cisa.gov/vulnerability-disclosure-policy-template

[37] Google, "Project Zero: Vulnerability Disclosure Policy." [Online]. Available: https://projectzero.google/vulnerability-disclosure-policy.html

[38] S. B. Merriam and E. J. Tisdell, *Qualitative Research: A Guide to Design and Implementation*, 5th ed. John Wiley & Sons, 2025.

[39] G. Guest, A. Bunce, and L. Johnson, "How Many Interviews Are Enough?" *Field Methods*, vol. 18, no. 1, pp. 59–82, 2 2006.

[40] V. Braun and V. Clarke, "One size fits all? What counts as quality practice in (reflexive) thematic analysis?" *Qualitative Research in Psychology*, vol. 18, no. 3, pp. 328–352, 7 2021.

[41] F. Li, L. Rogers, A. Mathur, N. Malkin, and M. Chetty, "Keepers of the Machines: Examining How System Administrators Manage Software Updates For Multiple Machines," in *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. Santa Clara, CA: USENIX Association, 8 2019, pp. 273–288. [Online]. Available: https://www.usenix.org/conference/soups2019/presentation/li

[42] S. Kraemer and P. Carayon, "Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists," *Applied Ergonomics*, vol. 38, no. 2, pp. 143–154, 3 2007.

[43] Gartner Inc., "Gartner Glossary - Small And Midsize Business."

[44] S. Frei, M. May, U. Fiedler, and B. Plattner, "Large-scale vulnerability analysis," in *Proceedings of the 2006 SIGCOMM workshop on Large-scale attack defense - LSAD '06*. New York, New York, USA: ACM Press, 2006, pp. 131–138.

[45] M. Shahzad, M. Z. Shafiq, and A. X. Liu, "A large scale exploratory analysis of software vulnerability life cycles," in *2012 34th International Conference on Software Engineering (ICSE)*. IEEE, 6 2012, pp. 771–781.

[46] N. Alexopoulos, M. Brack, J. P. Wagner, T. Grube, and M. Mühlhäuser, "How Long Do Vulnerabilities Live in the Code? A Large-Scale Empirical Measurement Study on FOSS Vulnerability Lifetimes," in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, 8 2022, pp. 359–376. [Online]. Available: https://www.usenix.org/conference/usenixsecurity22/presentation/alexopoulos

[47] N. Alexopoulos, A. Meneely, D. Arnouts, and M. Mühlhäuser, "Who are Vulnerability Reporters?" in *Proceedings of the 15th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*. New York, NY, USA: ACM, 10 2021, pp. 1–12.

A Interview Protocol

A.1 General questions:

1. Can you introduce yourself, share your job title, and describe your role within your organization?
2. How long have you worked at your current organization?
3. Does your organization have a CVD program? If so, since when?

A.2 Questions for vendors with a CVD program:

A.2.1 Pre-adoption considerations

1. What were the main goals your organization aimed to achieve by implementing the CVD program?
2. What pros and cons did your organization consider before adopting the CVD program?

A.2.2 Vulnerability discovery and remediation in practice

3. Can you walk me through the steps your organization takes when receiving a vulnerability report? Does this process differ from how internally discovered vulnerabilities are handled?
4. On average, how many vulnerabilities does your organization discover internally in its products each year?
5. On average, how many vulnerabilities are reported through the CVD program each year?
6. What channels does your organization use to inform customers about vulnerability fixes?

A.2.3 Interaction with vulnerability reporters

7. How would you generally describe your organization's interactions with vulnerability reporters?
8. Can you share an example of a positive interaction with ethical hackers, and explain what made the experience successful?
9. Can you share an example of a negative interaction with ethical hackers, and explain what made the experience challenging?
10. Have there been cases where researchers publicly disclosed vulnerabilities instead of reporting them through your program? How did you handle the situation?

A.2.4 Post-adoption challenges and countermeasures

11. What challenges has your organization encountered with the CVD program, and what measures have been taken to mitigate them?
12. Do the challenges and mitigations differ between vulnerabilities discovered internally and those reported through CVD programs? If so, how?

A.2.5 Recommendations for organizations considering CVD implementation

13. What advice would you give to organizations considering implementing a CVD program?
14. Is there anything else you would like to share that can be valuable for our research?

A.3 Questions for vendors without a CVD program:

A.3.1 Pre-adoption considerations

1. Has your organization internally discussed adopting a CVD program?
2. What pros and cons were considered?
3. What is your current perspective on CVD programs?

A.3.2 Vulnerability discovery and remediation in practice

4. Can you walk me through your organization's process for discovering and remediating vulnerabilities?
5. On average, how many vulnerabilities does your organization discover internally in its products each year?
6. Does your organization receive vulnerability reports from external researchers? If so, how many do you receive on average each year?
7. Does the process for handling vulnerabilities reported by external researchers differ from the process for internally discovered vulnerabilities? If so, how?
8. What channels does your organization use to inform customers about vulnerability fixes?

A.3.3 Interactions with vulnerability reporters

9. How would you generally describe your organization's interactions with ethical hackers?
10. Can you share an example of a positive interaction with ethical hackers, and explain what made the experience successful?
11. Can you share an example of a negative interaction with ethical hackers, and explain what made the experience challenging?
12. Have there been cases where researchers publicly disclosed vulnerabilities instead of reporting them to your organization? How did you handle the situation?

A.3.4 Post-adoption challenges and mitigations

13. Does your organization face any challenges with vulnerabilities reported by external researchers? How do these challenges compare to those of internally discovered vulnerabilities, and why?
14. What measures, if any, has your organization implemented to address the challenges of vulnerabilities reported by external researchers?

A.3.5 Recommendations for organizations considering CVD implementation

15. What advice would you give to vendors considering implementing a CVD program?
16. Is there anything else you would like to share that can be valuable for our research?