Computers & Security 123 (2022) 102936

Contents lists available at ScienceDirect

Computers & Security

journal homepage: www.elsevier.com/locate/cose

Coordinated Vulnerability Disclosure programme effectiveness: Issues and recommendations

T. Walshe*, A.C. Simpson

Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, United Kingdom

ARTICLE INFO

Article history:
Received 22 June 2022
Revised 22 September 2022
Accepted 25 September 2022
Available online 27 September 2022

Keywords:
Coordinated vulnerability discovery
Bug bounty programmes
Responsible disclosure programmes

ABSTRACT

Coordinated Vulnerability Disclosure (CVD) programmes leverage a global network of independent security researchers (hackers) to support pre- and post-deployment security. Organisations are increasingly adopting Bug Bounty Programmes (BBPs) and Vulnerability Disclosure Programmes (VDPs) to outsource work from internal security teams, and are able to utilise the results from a programme to help shape their Software Development Life Cycle (SDLC) processes. Motivated by the question *How effectively are organisations utilising CVD programmes?*, we aim to address two issues concerning the operation of CVD programmes. First, it is necessary to identify the pre- and post-launch issues faced by programme operators that inhibit effective operation. Second, organisations stand to benefit if they are able to use the results of a CVD programme outside of the typical reporting-triaging information flow between a hacker and the operator. As such, it is useful to explore how the results of a CVD programme influence change across the SDLCs of real-world organisations and measure the extent to which this occurs. We report upon the results of a qualitative study based on the outcomes of 39 survey responses and eight semi-structured interviews with individuals involved in the operation of CVD programmes. It is found that the fears and issues faced by organisations are similar to those identified in earlier studies, suggesting that there has been little development in preventing prevalent problems faced by CVD programme operators. High volumes of low-quality, low-value reports still burden operators and consume resources. It is also found that organisations use the information contained within vulnerability reports to influence change in a number of security activities, namely testing, communication processes, and the specification of security requirements. Finally, based on the responses from the surveys and interviews, we provide recommendations to those looking to establish a CVD programme.

1. Introduction

Many organisations choose to operate Coordinated Vulnerability Disclosure (CVD) programmes (Silomon et al., 2022), opening up their assets to large numbers of curious white-hat hackers (see Vandervelden et al. (2021) for a further discussion of the types of hackers) in the hope that these hackers will discover vulnerabilities and disclose the details to the security team (Cavusoglu et al., 2007). This can provide benefits to pre- and post-deployment security as vulnerable assets are patched (Takanen et al., 2018). The broad technical background of the global community of hackers has been responsible for identifying and reporting tens of thousands of vulnerabilities in a wide range of technologies and asset types (Votipka et al., 2018; Walshe and Simpson, 2020).

In exchange for the submission of a valid vulnerability report, hackers may be rewarded with monetary payouts, 'swag' (such as a company t-shirt), public recognition, or employment (Hata et al., 2017). Although hackers may be motivated by the eye-catching rewards offered as payouts from many Bug Bounty Programmes (BBPs) (Walshe and Simpson, 2022; Zhao et al., 2014), such as the $1,000,000 bounties offered by Apple (Hern, 2019), some act out of altruism and report vulnerabilities to ensure the privacy, safety and security of the general public (HackerOne, 2021).

In recent years, the number of organisations operating a CVD programme has continued to increase (Walshe and Simpson, 2020; Zhao et al., 2015). This is particularly evident on bug bounty platforms: third-party organisations that host CVD programmes, facilitating access to a pool of hackers (in some cases with ver-

* Corresponding author.
E-mail addresses: Thomas.Walshe@cs.ox.ac.uk (T. Walshe), Andrew.Simpson@cs.ox.ac.uk (A.C. Simpson).

https://doi.org/10.1016/j.cose.2022.102936

ified identities), global payment systems, and optional paid-for report verification and management services (Bugcrowd, 2022; HackerOne, 2022). The number of CVD programmes available on bug bounty platforms has grown year-on-year, and now includes major organisations, such as the U.S. Department of Defense, IBM, Uber, Atlassian, and Cloudinary (Bugcrowd, 2022; HackerOne, 2022). The acceptance and growth of this activity may be attributed, in part, to its incorporation into secure design frameworks (such as the Building Security in Maturity Model (BSIMM) (Synopsys, 2019)), professional standards (such as ISO/IEC 29147 (International Organization for Standardization, 2022)), and binding operational directives by the U.S. Department of Homeland Security (such as BOS 20-01 (U.S. Department of Homeland Security, 2022)).

Despite the widespread usage of CVD programmes and the increasing maturity of the security activity, there are still many barriers to adoption, including: large volumes of low quality reports (Al-Banna et al., 2018), lack of hacker motivation (Walshe and Simpson, 2022), high operating costs (Walshe and Simpson, 2020), and a general distrust of hackers (Follis and Fish, 2022). A qualitative study involving the operators of CVD programmes has the potential to both uncover the difficulties faced by operators during the pre- and post-launch phases and assess how effectively organisations are utilising their programmes. It is hoped that the results and insights presented in this paper can be used by organisations wanting to deploy their own CVD programme and also help inform those already operating one.

The structure of the remainder of the paper is as follows. Section 2 explores the background of CVD programmes, discusess related studies, and presents the motivation for the work described in this paper. The methodology is detailed in Section 3. The background of the respondents is presented in Section 4. Having introduced our three research questions in Section 2, we consider the results in Sections 5–7. Finally, in Section 8, we summarise the contributions of the paper and give consideration to areas of potential future work.

2. Background and motivation

In this section we give consideration to the background to, and the motivation for, our contribution. The background to the research is outlined in Section 2.1, which also provides context for the objectives of the study. A summary of related work is presented in Section 2.2. Finally, Section 2.3 presents the motivation for the study undertaken, as well as our research questions.

2.1. Background

Secure Software Development Lifecycles (SSDLs) represent an extension of the traditional Software Development Life Cycle (SDLC) development paradigm, seeking to better integrate security activities and checkpoints throughout the entire product lifecycle (Beckers et al., 2015; Jones and Rastogi, 2004). By considering security at all phases (i.e. during planning, development, and deployment), the use of an SSDL framework is intended to prevent the presence of architectural flaws and software and hardware vulnerabilities in the deployed product (Ardi et al., 2007). Aside from the obvious benefit to security, the mitigation of errors during earlier stages in the product lifecycle has been shown to decrease the total cost of development and future maintenance (Pressman, 2005).

Created in 2008, and first published in 2009, the BSIMM presents a data-driven approach to SSDL framework creation (Synopsys, 2019). The BSIMM has grown from measuring the security activities of nine organisations (BSIMM1, 2009) to using measurements from 128 major organisations (BSIMM12, 2021) (Synopsys, 2022). Its current iteration (BSIMM12) presents a collection of 122 curated security activities across the domains of: governance, intelligence, SSDL touchpoints, and development. Each domain is decomposed into practices, which are, in turn, decomposed into individual security activities.

There are three mature activities of particular interest to us within the Configuration Management and Vulnerability Management (CMVM) practice. First, CMVM 3.4 advocates for the operation of a bug bounty programme. The activity was first identified in BSIMM-V (2013) but not measured until BSIMM6 (2014) (Synopsys, 2022). It has seen steady growth within the BSIMM: from being present in 4% of measured organisations in BSIMM6 to 15.6% of organisations in BSIMM12 (Synopsys, 2022). Second, the new addition of CMVM 3.7 in the BSIMM promotes the streamlining of incoming responsible vulnerability disclosure reports, and encourages organisations to improve the ease-of-access of reporting information to hackers (such as visibility of security email addresses and security.txt documentation (Poteat and Li, 2021)). Third, CMVM 3.2 encourages a systematic refinement of the SSDL framework using the feedback from operations teams. While this security activity has been included in the BSIMM since inception, it has seemingly failed to gain popularity within organisations (Synopsys, 2022). Within the context of this paper, this activity is of particular interest. In order to fully utilise the information contained within vulnerability reports (and any metadata), it would be prudent for an organisation to analyse reporting patterns and make necessary changes to the SSDL / SDLC processes to prevent reoccurrence. The inclusion of CVD programmes within data-driven SSDL frameworks, together with recommendations relating to the use of data in a 'feedback-to-SSDL step', demonstrate the growing acceptance and utilisation of CVD programmes within modern organisations.

In this paper we define a CVD programme as the broad security activity that involves the disclosure of vulnerability information from an external white-hat hacker (individuals that 'positively' impact security through the disclosure of vulnerabilities (Silic and Lowry, 2021)) to an organisation (Householder et al., 2017). Bug bounty and vulnerability disclosure (often called responsible disclosure) programmes can be considered as two commonly employed types of CVD programme. For the purpose of this work we define a BBP as a CVD programme that explicitly offers monetary rewards — often outlined in the programme policy — for the submission of eligible vulnerability reports. In contrast, we define a Vulnerability Disclosure Programme (VDP) as a CVD programme that does not explicitly offer (nor imply) monetary rewards. Nevertheless, intangible rewards (e.g. listing in a hall of fame, reputation points, etc.) and 'swag' may still be rewarded to successful hackers. Furthermore, both types of programme may offer one-off monetary rewards (not defined by programme policy) at the discretion of the programme manager for particularly insightful or severe reports.

As highlighted by Formosa et al. (2021) and Manjikian (2017), BPPs may be ethically questionable if the operators simultaneously encourage participation but fail to provide adequate authorisation for the activities. In an effort to provide some legal assurances to hackers, organisations may choose to publish 'Safe Harbor' guidelines that permit 'good faith' vulnerability research. A 'Full Safe Harbor' (see the Disclose.io project for templates (Disclose.io, 2022)) is achieved if explicit authorisation is given to hackers conducting research an organisation's assets, and may help alleviate the risks that arise due to certain anti-hacking or anti-circumvention laws, (e.g. the Computer Fraud and Abuse Act (CFAA)) (Etcovitch and van der Merwe, 2018). Further discussion of the legal risks and implications of bug bounties can be found in Ellis and Stevens (2022).

As discussed in Section 1, many CVD programmes are hosted on bug bounty platforms, such as HackerOne[1] and Bugcrowd[2], and, despite the name, bug bounty platforms often allow organisations to host a VDP. These platforms act as centralised directories for programmes, allowing user-bases of hundreds-of-thousands of hackers to view pertinent information about an organisation's CVD policies and the extent to which hackers have participated in the past. As market participants, platforms help to resolve information asymmetries between programmes and hackers (Wachs, 2022; Walshe and Simpson, 2022).

2.2. Related work

There have, in recent years, been numerous qualitative studies exploring the backgrounds and the relationships between hackers, security professionals and organisations. For example, a study by Votipka et al. (2018) presents results from 25 semi-structured interviews in order to compare and contrast the experiences between internal security professionals ('testers') and hackers. Similarly to Hafiz and Fang (2016), the authors consider how individuals select areas to search, make use of automated tools, and communicate with other stakeholders. Although hackers typically work outside of a team structure, Cuevas et al. suggest that collaborative structures may benefit individuals when searching for vulnerabilities as double-work can be avoided (Cuevas et al., 2022). Witschey et al. (2015) also consider the adoption and usage of security tools amongst developers. The contributions of Votipka et al. (2018) and Hafiz and Fang (2016) both identify groups of hackers that search for vulnerabilities as a hobby. Respondents explain that this is done for personal enjoyment, but also as a good opportunity to learn via real-world examples (Votipka et al., 2018). In HackerOne's 2019 report (HackerOne, 2019), learning is cited as one of the key non-monetary motivating factors for hackers on the platform. This conflicts with the views of Malladi and Subramanian (2019) that monetary rewards act as the sole source of motivation for hackers to search for vulnerabilities.

An earlier study by Al-Banna et al. (2016) used 66 survey responses and 32 interviews to understand how the expertise of a hacker is indicated to other hackers and security professionals. The authors highlight the value to an organisation that can be found through the invitation of hackers to private programmes, such as those hosted on third-party platforms (e.g. HackerOne), through the identification of expertise indicators. Public programmes do not allow for an organisation to restrict access to low-skilled hackers. As such, without the ability to filter by expertise, a programme may receive an overwhelming majority of invalid reports ('noise'). Supply shocks due to the COVID-19 pandemic may have exacerbated such problems (Zrahia et al., 2022). Automated approaches to report validity checking, such the report classifiers by Fan et al. (2018), or by extending the topic models proposed by Xia et al. (2016), may be beneficial to organisations looking to filter out valid reports from the incoming noise. However, practitioners may be wary of automated report management systems in practice due to reliability (Zou et al., 2018). A study by Li and Zhao (2022) explores how operators many utilise governance mechanisms (such as definitions of in-scope assets) may be used to control the quantity of submissions to a programme.

Smith et al. (2020) explore the use and effectiveness of internal red teams (offensive security) at Microsoft using the results of 17 interviews with red-team members. Unlike hackers, internal red teams may utilise insider knowledge and shortcuts (such as internal network access) to more quickly identify vulnerabilities or consider more advanced attack scenarios (Hilton et al., 2017). Alomar et al. (2020) also use interviews to explore vulnerability discovery and management from the perspectives of internal testing teams, penetration testers, red teams, blue teams (defensive security) and purple teams (offensive and defensive), and through the employment of BBPs. The authors highlight the need for a mature security culture, effective internal communication channels, and acceptance by managers before the aforementioned security teams can be fully utilised.

Using the results of 36 interviews, Al-Banna et al. (2018) investigated the organisation perspective on the usage of CVD programmes. The authors uncovered the fears that many organisations express towards the use of CVD programmes. Of the interviewees, 10 were actively involved in the operation of a CVD programme and provide insights as to whether these fears are realised in practice. In addition, ongoing issues during programme operation, as well as countermeasures, are presented. Many of the fears and issues presented in the study are echoed in other qualitative and quantitative work. A distrust of hackers by some organisations is also found by Tanczer (2020). A lack of hacker motivation and a difficulty in maintaining participants is reported by Walshe and Simpson (2022).

2.3. Motivation

The focus of the study described in this paper is, in part, inspired by the creation and use of a Bug Bounty Lifecycle (BBLC) framework by Media (2022). The ambition of the operator is to maximise the value derived from the BBP using a lifecycle approach that incorporates the findings from hackers to help guide organisational change in respect to security efforts and activities (Media, 2022). Although there have been many studies that consider aspects surrounding the usage of CVD programmes by organisations, there has been rather less investigation into the extent to which organisations use the results of a CVD programme to influence change in their SDLCs.

Furthermore, the continued adoption of CVD programmes provides motivation to explore how organisational perspectives have changed over time. Motivated by the work of Al-Banna et al. (2018) (who conducted interviews between 2015 and 2016), it is useful to explore how the pre-launch fears and post-launch issues have changed as the usage of CVD programmes has matured in recent years. It is also useful to assess the recommendations of Alomar et al. (2020) in the context of BBPs and VDPs.

It is hoped that, by identifying how organisations use the findings from their programmes to bring about change, as well as how they overcome fears and issues that arise, the work described in this paper will inspire other organisations to take onboard the recommendations put forward in this study and more effectively utilise their CVD programmes.

We consider *How effectively are organisations utilising CVD programmes?* through answering the following research questions:

1. What are the current issues and fears faced by programme operators pre- and post-launch?
2. To what extent do organisations use the findings from their CVD programme to direct change throughout their SDLC?
3. How should organisations looking to leverage the work of hackers best set up a CVD programme?

Question 1 seeks to re-examine the findings of Al-Banna et al. (2018) using new data collected up to five years after their original study. If the same fears and issues persist, it perhaps demonstrates the existence of systemic flaws in the current methods to promote and inform organisations about the usage of CVD programmes. The aim of Question 2 is to identify

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

the extent to which organisations are currently using the findings from hackers to influence changes within the organisation, helping to improve security. Furthermore, by identifying which activities are being influenced by the findings, it is hoped that more organisations may be inspired to reflect upon their own processes. In conjunction with the recommendations put forward by Alomar et al. Alomar et al. (2020), Question 3 considers how best an organisation should set up a new CVD programme. In contrast to the earlier contribution of Alomar et al. Alomar et al. (2020), this study also considers the use of VDPs in addition to BBPs.

3. Methodology

The process of selecting candidates for the study is discussed in Section 3.1. The survey design methodology is outlined in Section 3.2 alongside the survey distribution process. A brief overview of the interview process is provided in Section 3.3. Finally, the data analysis methods are presented in Section 3.4.

3.1. Candidate selection

Organisations operating BBPs and VDPs were considered, as both are categorised as subtypes of CVD programmes (see Section 2.1). Organisations were selected based on their inclusion on a bug bounty platform (HackerOne (2022), Bugcrowd (2022), iNTiGRiTi (2022), Cobalt (2022), and YogOsha (2022)), listed in the Disclose.io (2022), included in the *Awesome Bug Bounty* GitHub repository (Joshi, 2022), listed on FireBounty (2022), mentioned on the Google Play Security Reward Program (GPSRP) (Alphabet, 2022), or identified from any additional public mentions of a CVD programme on a company website.

Many of the bug bounty platforms provide the names of organisations operating a VDP, as well as those operating a BBP. This resulted in 950 unique organisations with a corresponding contact email address; manual searching and matching was required in many instances to identify a suitable line of contact for organisations listed without an email address. A large number of candidate organisations were selected to help the generalisability of the results, but also to mitigate the poor response rates that are typical with web-based surveys that recruit via email (Manfreda et al., 2008; Sappleton and Lourenço, 2016).

3.2. Survey design and distribution

Based on the recommendations found in survey design literature (Fink, 2003; Krosnick, 2018; Lietz, 2010; Sanchez, 1992), a typical iterative design process was used to construct and design questions (and groups of questions) in order to improve readability, ensure good narrative flow, and avoid bias. Respondents were presented with a total of 24 open-ended and closed-form questions; branching was employed to prevent respondents from being presented with questions that would have been irrelevant to their organisation. Four of the questions related to obtaining consent from the respondents and the option to participate in a follow-up interview. Considerations were made to ensure that the total time to complete the survey did not exceed the approximate duration of 15–20 minutes, as reducing the burden on the participants can help improve completion rates and may also impact the quality of the collected open-ended data (Galesic and Bosnjak, 2009).

Prior to distribution, care was taken to ensure that the study complied with commonplace professional and ethical codes of conduct, such as the ACM Code of Ethics (ACM, 2022). A study by Buchanan and Hvizdak (2009) found that issues around consent, confidentiality, security and data storage (among others) were of particular importance to web-based surveys and survey tools. As such, additional refinement to the survey was made to minimise any personal information collected about a participant. To comply with departmental requirements, the Jisc Online Surveys platform[3] was used throughout. Ethics approval was sought from, and granted by, the University of Oxford's Computer Science Departmental Research Ethics Committee (CS-DREC)[4].

Organisations were contacted between February 2021 and October 2021. From the 950 organisations contacted, 278 individuals accessed the survey site, leading to 39 complete responses. Individuals were not compensated for either the survey or follow-up interviews. On average, approximately 6% of emails sent were undeliverable to the target organisation, suggesting that some of the information contained on public bounty lists may be outdated.

3.3. Interview process

The final section of the survey allowed participants to express an interest in participating in a follow-up interview and give consent to being contacted. From the 39 survey respondents, 14 expressed an interest in a follow-up interview, leading to 8 interviews being conducted. Semi-structured interviews were subsequently held on the Microsoft Teams platform. Information provided in the survey response and public information about the organisation's CVD programme was used to guide the discussions.

3.4. Data analysis

The responses to all closed-form questions were separated from the survey data. Many of these questions were followed up by open-form questions, as suggested by Schuman and Presser (1979). However, on their own, the categorical data were analysed to produce insights into sentiment and programme structure. Despite the use of semantically ordinal scales (strongly agree, agree, neither agree nor disagree,...) for the reporting of a respondent's sentiment towards a question, previous studies have highlighted the suitability of scales for use in parametric statistical analysis (Desselle, 2005; Kerlinger, 1966). As such, statistics from the Likert scales are presented where meaningful (Sullivan and Artino Jr, 2013). A 'sentiment' score is computed using the arithmetic mean from the Likert scales following numerical assignment (e.g. strongly agree = 2, agree = 1, neither agree nor disagree = 0,...).

The survey data and interview transcripts were analysed using the thematic analysis technique described by Braun and Clarke (2006). Their qualitative techniques are widely used in relevant computer science literature, particularly when applied to the results of surveys and interviews (Al-Banna et al., 2018; Emami-Naeini et al., 2020; Kapoor et al., 2022; Smith et al., 2020; Tanczer, 2020). As such, their approach was deemed to be appropriate for this study.

The process is as follows. First, all relevant data is centralised and read over, and initial observations are noted. Then, interesting features within the data are systematically coded and the relevant data is extracted and grouped. Care is taken to preserve the context of the data by including surrounding sentences in the extract (Bryman, 2016). The codes (and associated data) are then iteratively grouped and refined to ensure that the identified themes are consistent with a given group of codes, and across the entire corpus. The themes resulting from the application of this process were used to help address the research questions presented in Section 2.3.

[3] https://www.onlinesurveys.ac.uk

[4] Approval number CS_C1A_20_029 and approval date 9th of December 2020

4. Background statistics

Overall, the survey had 39 responses leading to 8 follow-up interviews. Of the survey respondents, 16 operated a VDP, 12 operated a BBP, 10 operated both, and one respondent noted that their programme fit neither designation. From the operators of a BBP or both type of programme, 12 used a platform (primarily HackerOne and Bugcrowd) and 10 did not. 31 respondents offered rewards and eight (all VDPs) offered none. Full Safe Harbor was the most common form of protection offered amongst participants (14), however a large proportion of respondents offered none (8), or were unsure (8).

Of the interview respondents, three operated a VDP, three operated a BBP, and two operated both types of programme.

5. Issues and fears surrounding CVD operation

We consider the first of our research questions in this section. Fears expressed by participants prior to the launch of a programme are presented in Section 5.1. This is followed by a discussion of the issues faced by operators after the launch in Section 5.2. The countermeasures enacted by participants to combat the aforementioned issues are displayed in Section 5.3.

5.1. Pre-launch fears

From the results of the survey and interviews, five themes relating to organisational pre-launch fears are identified: use of hackers (Section 5.1.1), distrust of hackers (Section 5.1.2), communicating with hackers (Section 5.1.3), volume of reports (Section 5.1.4), and internal limitations (Section 5.1.5).

5.1.1. Fear: Usage of hackers

There is concern about the poor quality of reports submitted by hackers. Understandably, duplicate and invalid reports are of little-to-no value to an organisation. Furthermore, a lack of skill or a lack of incentive may result in the submission of shallow or low quality reports. One participant mentioned that they had a fear of a "high volume of low value and duplicate reports (this turned out to be a very valid fear)" [P3]. This sentiment is explicitly echoed by many of the participants.

5.1.2. Fear: Distrust of hackers

Despite the growing adoption of CVD programmes and widespread stories of the successes of hackers in finding vulnerabilities, there exists a distrust of hackers within some organisations. From the perspective of some departments, the unorthodoxy of allowing external and unknown hackers to explore an organisation's assets induces fear. Three of the participants note that their legal teams raised early concerns and issued guidance prior to launch: "Our legal team raised concerns about giving access to external people" [P30]. There is a fear that, following a negative interaction with a hacker, retribution may be taken in the form of full disclosure or through reputation damage on social media. Both [P26] ("we'd get something wrong and be dragged on social media") and [P11] ("there is also the fear that a reporter will talk badly about you on Twitter or similar") expressed this fear.

5.1.3. Fear: Communicating with hackers

For organisations opening up lines of communication with hackers, there are two areas of concern: disagreement and harassment. Participant [P19] expressed an early fear that hackers may not agree with the re-assessment of the severity of a vulnerability: "they would find a problem that we would not feel is not severe and the hacker would disagree that the issue is severe and there would be a conflict of interests". Additionally, [P20] noted a "fear of dealing with harassment from impatient researchers as we got our processes in order". Within the early stages of programme operation, an organisation may not have the processes in place to respond to incoming reports in a time-frame that is preferable to all hackers. In such a scenario, "chasing by reporters" [P32] may only impede that timely operation of a programme.

5.1.4. Fear: Volume of reports

The uncertainty around the number of incoming reports is another fear held by organisations. Primarily, organisations fear a volume of reports that is too large to deal with effectively: "could our systems and processes cope with an influx of disclosures at the beginning?" [P9].

5.1.5. Fear: Internal limitations

A number of internal limitations could impede the successful operation of a new programme and induce anxiety in those involved. There are three sub-themes that relate to the fears participants held in relation to internal limitations.

First, an organisation's employees may not posses the skills required to operate a programme correctly or fix the identified vulnerabilities. The inability to fix an exposed vulnerability not only puts an organisation at further risk to unknown hackers, but may sour communications with the original reporter, potentially leading to full disclosure. In summary, "you launch a programme but you don't have the expertise to evaluate the submission that you receive, you are going to have a bad time!" [P30].

Second, a failure to address incoming report in a timely manner is identified as a fear. Not only do longer response times displease hackers, but the burden of responding to reports has the potential to overwhelm a team: "we also had to discuss how much time triaging reports would take" [P15].

Third, programme operators fear that their budget, or lack thereof, may inhibit the success of an upcoming programme. Many attribute this (prior to launch) to low or no bounty payouts as a reward for successfully finding a vulnerability, with the major constraint being internal budget allocation. "Our fear, as a not-for-profit organisation, was the limited impact of a CVD programme (e.g. low quality vulnerability reports) if you do not award large monetary rewards. Which proved to be true" [P34].

5.2. Post-launch issues

Three themes are identified when considering the issues faced by organisations post-launch: issues with reports (Section 5.2.1), issues concerning the behaviour of hackers (Section 5.2.2), and internal issues (Section 5.2.3).

5.2.1. Issue: Reports

There are four issues surrounding the vulnerability reports submitted by hackers: quality, volume, scope, and severity. These are discussed in turn.

Respondents commonly cited the poor quality of incoming reports as a major issue facing their CVD programme, with many reports being: false positives (invalid), inconsequential (valid, but not useful), duplicates, incomplete, or relevant to another organisation. Aside from the validity of the information contained within, many reports are written in bad English, fail to demonstrate the issue, or are in an unreadable file format. The use of automated scanning tools, particularly by low skilled hackers [P9], often leads to a significant volume of low quality results: "it seems like everyone who can run Burp cut-n-pastes their scanner output into our vulnerability report form" [P39].

Interestingly, respondents face differing issues when it comes to the volume of reports they receive. For some organisations, high volume goes hand-in-hand with poor quality: "we get a massive

amount of unhelpful reports" [P18]. Conversely, organisations may experience a drought of reports, perhaps due to unmotivated hackers, diminishing the returns from the programme.

A commonly observed problem is hackers ignoring (or in some cases arguing against [P28]) the programme scope set out by an organisation in their CVD policy. It is believed that some individuals fail to read the policies outlined by an organisation, instead choosing to test in both in- and out-of-scope assets. "Many researchers do not read (or understand) our disclosure policy, particularly with regard to in-scope systems" [P9].

The final issue concerning reports pertains to the over-inflation of the report severity. The cause of the issue may be twofold: potential monetary rewards are directly linked to the severity of a vulnerability (misaligned incentives), or, as an outsider, a hacker is unable to appropriately quantify the severity of a vulnerability in the context of an organisation's assets (information asymmetry or lack of hacker skill). As such, severity over-inflation is common, requiring reassessment by a skilled programme operator [P19].

5.2.2. Issue: Behaviour of hackers

There are four issues in connection with the behaviour of hackers: communication, motivation, rewards, and undesirable behaviour. These are discussed in turn.

For many organisations, the fear of poor communication with hackers has manifested as a prevalent issue during operation. Respondents reported that frequently the most 'vocal' hackers were those that had submitted an invalid report: "communication is generally poor, a lot of them refuse, don't understand that we are communicating, don't understand why what they have found is not a real issue" [P26].

The success of a programme is, in part, dependent on attracting high quality researchers. Once the 'low-hanging fruit' have been reported, subsequent discoveries require the continued effort of motivated hackers. For operators, enticing hackers back to a programme is problematic: "many researchers they're really focussed on low-hanging fruits" [P1]. An operator of a mature programme expressed a difficulty in motivating hackers to search for vulnerabilities in their assets as they are not trivial to test [P1]. Another respondent noted that previous reputation damage in the hacker community has dissuaded previous participants from returning [P20].

An organisation may carefully construct their policy documents to explicitly define client (customer) systems as being out-of-scope, especially if they not have ownership of the systems. One respondent ([P26]) noted that hackers have, in the past, tested against and incidentally exploited the vulnerabilities in client systems when searching for and then later reporting issues. Such behaviour is clearly extremely undesirable, and often cannot easily be prevented. Another respondent stated that, "as we do not offer bounties, we are unable to make some of them reveal their true identity and hence there is no legal recourse we may pursue in such cases" [P32].

5.2.3. Issue: Internal to organisations

Four issues are identified surrounding the internal issues faced by organisations post-launch: management, burden, pushback, and lack of skilled staff. These are discussed in turn.

Separation between the team responsible for managing the CVD programme and the teams responsible for developing or securing various company assets (while often necessary) causes management issues within some organisations. For large organisations, the process of identifying the responsible product owner, assigning the security issue, and managing communication can be a complex and time-consuming process [P3, P11, P21]. And, although many organisations make use of ticketing systems to resolve the aforementioned issues, management and tracking can still be challenging [P3, P34].

Despite the successful discovery, triage and internal communication of a vulnerability, several organisations report pushback from the teams responsible for a fix; a reluctance to fix valid security issues undermines the work of the operators and diminishes the benefits gained from running a CVD programme. Instead, respondents noted a "lack of willingness to prioritise remediation by some product teams (substantial issue for easily discoverable issues that are then repeatedly reported)" [P3]. However, when constrained by the choice of spending time on product development versus fixing low-severity vulnerabilities [P21], it is understandable that some overstretched teams may de-prioritise security: "We have not been internally diligent about following up on low-impact reports. While this is in many cases an issue of engineering resource allocation, multi-month turnarounds lead to researcher dissatisfaction and disillusionment" [P39]. Ultimately, for any organisation with security as a top priority, resisting pushback is necessary: "we had to use diplomacy to make them understand security wasn't optional and that it was a mandatory requirement" [P13].

An issue within some organisations is the absence of employees capable of tackling the demanding reports submitted by hackers. Furthermore, a lack of experience may fail to address broader issues surrounding the existence of a vulnerability: "junior engineers often can't properly assess the depth of issues or may patch a bug without addressing a systematic issue that could re-occur in other areas of code" [P37]. In one organisation, the triaging team for the CVD programme contained employees without any cybersecurity skills, potentially inhibiting the process of raising the alarm upon the submission of serious issues [P13].

5.3. Countermeasures

To counteract many of the fears and issues outlined earlier, a range of countermeasures may be employed. From these, there are four themes identified: internal changes (Section 5.3.1), use of platforms (Section 5.3.2), policy alterations (Section 5.3.3), and use of additional tools (Section 5.3.4).

5.3.1. Countermeasure: Internal changes

As highlighted in Section 5.2, myriad internal issues hamper efforts to operate a CVD programme following launch. For some organisations, a change in attitude and refinement of internal processes eased the burden on the operators and developers: "[we] obtained senior management support for prioritizing resolving vulnerabilities over developing new features" [P21]. In order to stem the tide of low-severity reports, one respondent altered internal policies by "treating certain lower-severity / theoretical issues as proper vulnerabilities, mostly to avoid getting repeated reports about them" [P26]. The lack of actionable countermeasures within this theme may highlight the continued difficulty in instituting internal change to counteract the issues previously raised.

5.3.2. Countermeasure: Use of platforms

To mitigate the drain on internal resources due to triaging and communicating with hackers, several organisations chose to make use of the 'managed programme' service offered by platforms such as HackerOne and Bugcrowd: "We have launched the program on [platform] so that this is no longer our internal ticketing triaging team responsible for triaging bug reports" [P34]. Furthermore, alongside the use of a managed programme, one respondent noted that private programmes, when operated on a platform, allow the operator to carefully invite hackers with verified identities, helping to partially resolve some of the fears of distrust and issues concerning hacker behaviour: "we accept only vetted users into our

program and have full control on their accounts ... so, not only are we accepting verified researchers but we also control their access, because if we want to revoke a credential we need the power to do it" [P30].

5.3.3. Countermeasure: Policy alterations

Alterations to programme policy were the most common type of countermeasure employed by respondents, resulting in two subthemes: reward modification and scope alteration. Altering the programme reward policy was used by organisations to both disincentivise the submission of undesirable reports (e.g. those of low severity vulnerabilities) and provide a greater incentive for hackers to search for vulnerabilities. In an attempt to reduce the number of lower severity reports, several respondents changed their bounty structures to offer no payouts for low severity reports: "we ultimately decided to only award bounties for medium- severity issues and above" [P15]. For those not in a position to offer monetary payouts, other tangible rewards proved successful: "we are now offering some 'swag' as rewards, which has led to a higher number of useful reports" [P34].

Using a well-defined scope allowed some respondents greater control of the types of reports that could be ignored or rejected during triage, thus helping to reduce the burden on the operators, with one respondent having "amended [the] scope several times to exclude low value classes of vulnerability" [P3]. By only considering worthwhile vulnerabilities in higher priority assets, some operators were able to more carefully control the running costs of the programme by limiting the number of incoming reports eligible for a bounty payout: "[we] had a pretty thorough program scope, so that cut down on payouts where we didn't think it was appropriate to reward" [P26].

5.3.4. Countermeasure: Use of additional tools

One operator of a private programme argued that the requirement for hackers to be authenticated while accessing assets allowed for the monitoring of hacker activity to be automated, giving an indication as to the effort being put into their bug bounty programme by hackers: "I set up a way to monitor how many researchers, more or less, are working on a weekly basis on our programme. You definitely see a trend that is going down over time" [P30].

5.4. Analysis

Inspired by relational visualisation between fears, issues and countermeasures by Al-Banna et al. Al-Banna et al. (2018), a diagram similar in concept is shown in Fig. 1 using the themes discussed in this section. In contrast to the visualisation of Al-Banna et al. Al-Banna et al. (2018), it is more logical to show how fears evolve into operation issues, and how they are subsequently affected by the enacted countermeasure.

5.4.1. Analysis of pre-launch fears

As shown in Fig. 1, 11 predominant fears were reported by respondents, covering the four themes described in Section 5.1. A comparison with the fears identified by Al-Banna et al. Al-Banna et al. (2018) reveals that many still persist within organisations pre-launch. And, as demonstrated by the issues faced by operators (see Section 5.2), the fears are often valid.

5.4.2. Analysis of post-launch issues

In common with Al-Banna et al. Al-Banna et al. (2018), issues surrounding the quality of reports, lack of hacker motivation, and burden on staff are still present to those currently operating CVD programmes. Despite research aimed at addressing the prevalence of low quality reports Ahmed and Lee (2020); Laszka et al. (2016); Malladi and Subramanian (2019); O'Hare and Shepherd (2022), it would appear that many organisations have yet to make use of the findings.

Table 1
Responses to "Have the results from the CVD programme prompted change to security practices?" (N = 38). The response from one organisation has been excluded as no answer was provided.

Programme Type	Yes	No	Unsure
BBP	9	3	0
Both	7	2	1
VDP	6	8	2
All	22	13	3

5.4.3. Analysis of suggested countermeasures

For most of the issues raised in Section 5.2, respondents offered suggestions of countermeasures that helped alleviate these ongoing issues. This is shown by the connections between issues and countermeasures in Fig. 1.

Comparing the suggested countermeasures with those of Al-Banna et al. Al-Banna et al. (2018) reveals many similar approaches used by organisations. Notably, 'third-party support' and 'limiting participants' can be realised through the adoption of managed and invite-only programmes. However, as we shall discuss in Section 7.2.4, these countermeasures often come at great monetary expense, so are not applicable to many smaller organisations. The absence of other countermeasures that address the need for verifiable hacker identities and those that help reduced the burden on operators exhibits a serious problem for the operators of many CVD programmes.

6. Impact on the SDLC

Inspired by the extent to which Verizon Media aims to integrate the results from their CVD programme into changes in the wider SDLC framework (aptly named the BBLC (Media, 2022)), respondents (N = 39) were asked to describe how they utilised report information in the context of organisational change. Results pertaining to the closed-form survey answers are presented in Section 6.1, and the open-ended survey and interview answers are presented in Section 6.2. Analysis and discussion is provided in Section 6.3.

6.1. Closed-form results

6.1.1. Impact on security practices

As shown in Table 1, a majority of organisations (22) have been prompted to make changes to security activities due to the results from the CVD programme. For 13 of surveyed organisations, the results of the programme did not elicit a change in practices, and a further three were unsure (and one participant did not provide an answer). There is a significant disparity between those exclusively operating a VDP versus those operating a BBP or both types of programme, with a half of those exclusively operating a VDP seeing little change in security practices owing to the programme's results.

6.1.2. Affected security activities

Respondents provided an indication of the security activities affected by the results of the CVD programme via a closed-form question, and the option to provide further information via an open-ended question. The results from the closed-form question are shown in Table 2. For those organisations that made changes, the average number of activities changed is somewhat similar across programme type: for BBP it is 3.3 changes; for Both it is 2.6 changes; for VDP it is 2.3 changes.

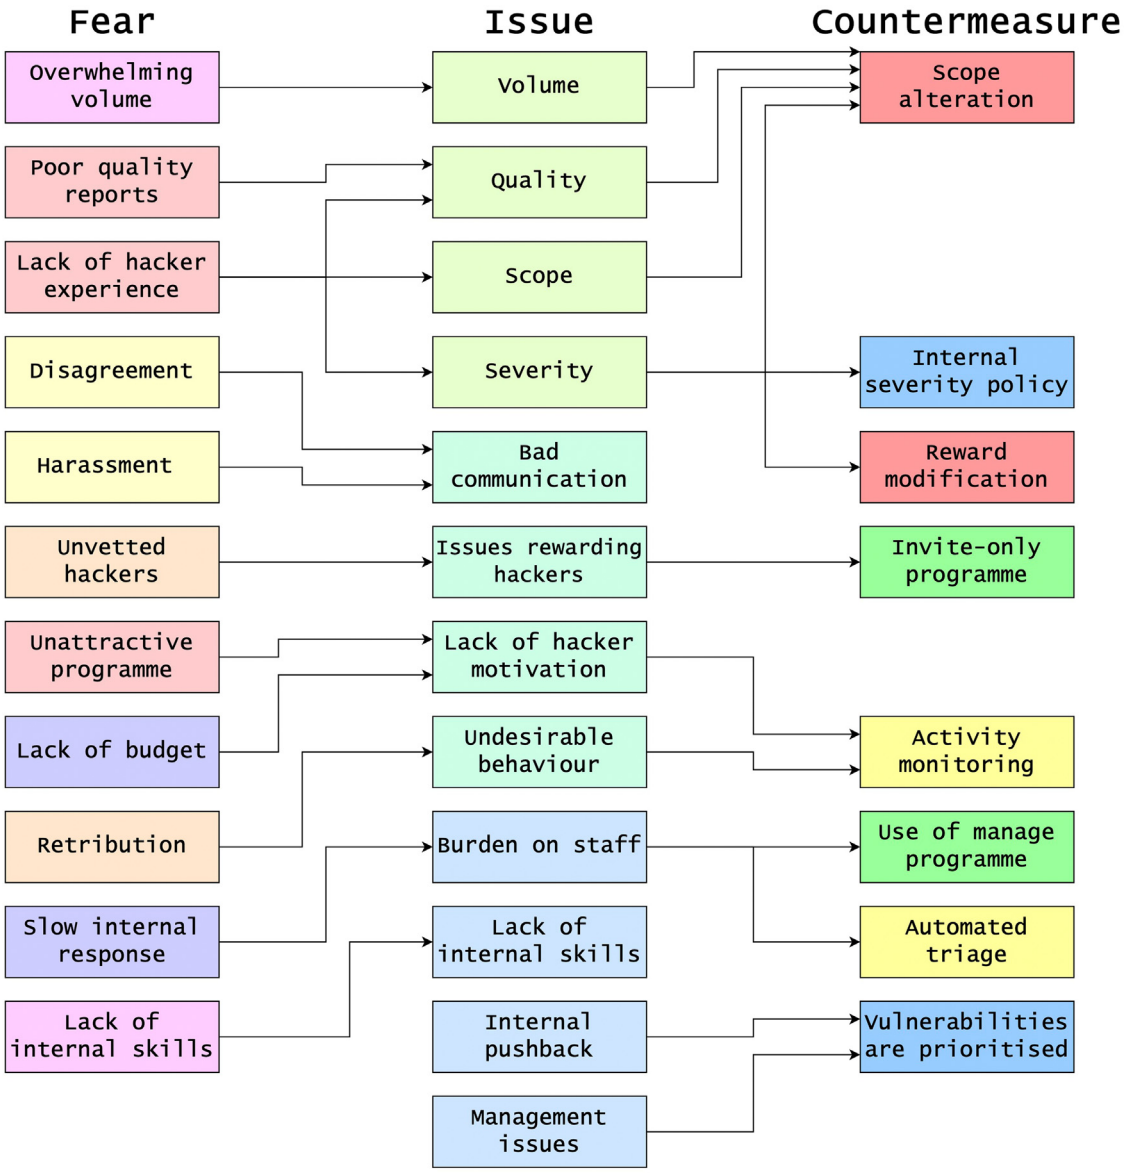

Fig. 1. Relationship between identified fears, issues, and countermeasures. Some connections have been hidden for clarity. Colour groups within a column represent items within the same theme (no relationship across columns).

Table 2
Number of respondents that made changes to particular security activities, given that their CVD programme results prompted change (N = 38). The response from one organisation has been excluded as no answer was provided.

Security activity	BBP	Both	VDP	All
Training for engineers	3	2	3	8
Focus or use of attack models	2	2	0	4
Specification of security requirements	6	3	1	10
Security standards and compliance measures	3	0	2	5
Communication between internal teams	6	4	2	12
Code review methods	1	2	1	4
Security testing tools or methods	7	2	4	13
Use of internal or professional penetration testing	2	3	1	6

6.1.3. Elimination of vulnerabilities

Organisations were also asked to comment on the effectiveness of their current approach to SDLCs and usage of CVD programmes in the context of vulnerability reappearance. If an effective process is in place, certain vulnerabilities may be less likely to reoccur once found and patched. Answers on a Likert scale from 'Strongly agree' to 'Strongly disagree' (see Section 3.2) are given ordinal representations (see Table 3). Those operating a BBP or both a BBP and a VDP expressed, on average, a slightly positive sentiment in relation to their ability to prevent vulnerabilities from reappearing (0.5 and 0.8, respectively). Conversely, the operators of VDP programmes expressed, on average, slightly negative sentiment (-0.5).

Table 3
Responses to "Certain categories of bug once found do not reappear due to changes in development and testing." (N = 39).

Response	BBP	Both	VDP	Other	All
Strongly agree (2)	2	1	0	0	3
Agree (1)	4	6	1	0	11
Neither agree or disagree (0)	4	3	8	0	15
Disagree (-1)	2	0	5	0	7
Strongly disagree (-2)	0	0	2	1	3

The sole respondent in the 'Other' category denoted a strong negative sentiment (-2). Across all respondents, the average sentiment was very slightly positive (0.10).

6.2. Open-ended results

Four themes are identified from the open-ended answers: improved employee training and education (Section 6.2.1), usage of internal security champions (Section 6.2.2), testing techniques (Section 6.2.3), and organisational strategy changes (Section 6.2.4).

6.2.1. Training and education

The ongoing arms race between increasingly sophisticated hackers and security professionals can leave developers without the requisite knowledge to address new and complex vulnerabilities that may affect their organisation (see Section 5.2). Furthermore, labour supply shortages of well-trained security professionals (Crumpler and Lewis, 2019; Furnell et al., 2017), and a general under-confidence in the cybersecurity-related proficiencies of new graduate developers (Catota et al., 2019; Conklin et al., 2014), result in many organisations around the world having development teams that lack the technical expertise to secure new and existing assets. It is therefore not surprising that both the closed-form and open-ended answers highlight the role of additional employee training and education following feedback from the CVD programme.

For some respondents, training serves as a reactive measure to mitigate the risks of repeated known attacks: "education sessions are held in response to patterns of reports or novel attacks" [P9] and "if we are starting to see a lot of vulnerabilities in one area, we may increase training on a selected training on a selected portion of the population" [P5]. One respondent described the role of internal workshops and presentations thus: "there has been some level of awareness work out of the back of the responsible disclosure programme which may help people think more about certain considerations" [P3]. They went on to say: "we have training paths now, focussing on certain commonly occurring vulnerability reports, so that more of the SOC team are able to triage reports when they come in" [P3].

6.2.2. Security champions

Security champions are individuals within an organisation, typically in development or engineering roles, that have a strong working knowledge of security topics, and may help guide or manage security activities Synopsys (2022). Respondents noted that an apparent shortfall of skills in certain teams could be partially mitigated by leveraging an internal network of security champions. This in some cases necessitated an expansion of the "security champions forums" [P4]. Strengthening and using the champions network in one organisation allowed more experienced security professionals to "jump [in] and and help validate more difficult to verify queries" during triage [P3].

6.2.3. Testing techniques

Automated testing techniques, if configured correctly, may help developers catch mistakes during development or configuration issues during deployment. Some respondents noted that a CVD programme's results prompted the organisation to revise or expand the testing methodologies currently employed: "we built some new detection for subdomain takeover problems directly as a result of reports" [P3].

6.2.4. Organisational strategy

Without periodic reviews, programme managers may be unable to see the prevalence of certain patterns of reports. Furthermore, if programme managers are not able to communicate the patterns they discover, systemic issues may persist, leading to a drain of resources as the CVD programme sees repeated reports on previously identified issues. These concerns prompted two respondents to better integrate the CVD programme outputs into a periodic and systematic review process: "we've got meetings that we have weekly where we discuss some of the disclosures that we are seeing to see if there is some strategy that needs to be changed. And, in some cases budget allocated for programme work (training, etc.)" [P5]. Another respondent noted that a formal review process was introduced with development teams to help effectively communicate current issues: "It's an involuntary meeting once a quarter, to look at the most recent security alerts. Here are the things that happened this quarter, it is just an educational process. Here is what is going on in the industry, bring your knowledge up to speed." [P33]

6.3. Analysis

As shown by the answers to the closed-form questions (Section 6.1) and open-ended questions (Section 6.2), a majority of the respondents noted that they had used the results from their CVD programme to help guide changes in security activities within their organisation. The use of testing tools and methodologies, alterations to internal communication processes, and specification of security requirements were the activities most often affected across all respondents.

Despite this, it is apparent from the results, particularly the interviews, that none of the respondents integrated the results of their CVD programmes to the extent of Verizon Media's BBLC. Instead, the utilisation is on a shallower level that suggests a more reactive approach to altering security activities based on the incoming reports. It is perhaps due to a lack of resources or a lack of internal integration that greater institutional change is not made. As noted in Section 6.2.4, few respondents shared details of the internal periodic information flows of report details that were used to help guide security strategies.

In summary, although organisations use the results from their CVD programmes to direct changes in security activities, at present this appears to be somewhat shallow within the organisations included in this study.

7. Recommendations to organisations

In order to aid those looking to set up a programme, respondents were asked about their set-up process and provided advice. Results for eight closed-form questions are presented in Section 7.1, and the results for the open-ended questions and survey responses in Section 7.2. Finally, analysis and discussion can be found in Section 7.3.

7.1. Closed-form results

In the survey, Questions 7–11 aimed to uncover the areas of difficulty encountered by programmes operators prior to or at the start of operation. Questions 20–22 asked respondents to reflect

Table 4
Average sentiment expressed across programme types for the closed-form survey questions (N = 38). "Other" programme type omitted due to non-response.

Question or statement	Average sentiment for type			
	BBP	Both	VDP	All
The process of launching a programme and the resources required were well understood	0.84	0.9	0.44	0.67
A well defined process was in place allowing for reports form hackers to be validated	0.49	1.4	0.44	0.67
A well defined process was in place to notify the responsible team about the presence of a vulnerability	1	1.3	0.75	0.96
Systems allowed for clear and easy communication between the hackers, the CVD team and the fixers of the bugs	0.83	1.3	0.63	0.83
There was a clear method used to set bounty levels or rewards that were appropriate for the organisation	0.34	0.9	0.31	0.43
The use of a CVD programme has benefited the security of assets	1.49	1.6	1.2	1.35
The operation of a CVD programme is a good use of resources	1.26	1.5	0.7	1.05
Would you recommend the use of a CVD programme to other organisations?	1.51	1.6	1.32	1.41

upon their overall experience operating a CVD programme. Summaries of the average sentiments for each are shown in Table 4. For each question, significance testing of the average sentiment is performed at the 5% level with Welch's T-test (unequal variance assumed), using the Bonferroni correction where appropriate, to compare the effects of: programme type, usage of a platform (for BBP and Both), and management style (fully- or semi-managed programmes compared to independent management on a platform). Unless significant results are obtained and presented, it can be assumed that no significant difference in the average sentiment is present at the 5% level.

As shown in Table 4, although the average sentiment for setting bounty levels and rewards is positive (0.43), many respondents expressed uncertainty. Comparing the sentiment of those using a managed programme (1.42) versus an independent programme (0.38) on a platform reveals a significant difference ($p = .0496$) at 5%, suggesting that platform managers are able to clearly convey their rationale behind setting initial reward structures. As explained by [P30], this typically involves comparing factors such as company size, industry, level of security maturity and budget to the reward structures of existing public and private programmes on a platform.

The final question in the survey asks organisations to provide an indication as to their sentiment when recommending the use of a CVD programme to other organisations (see Table 4). The overall sentiment is largely positive towards the use of CVD programmes (1.41). When tested at the 5% level, there is a significant difference in the average sentiment expressed by those using managed programmes (2.0) and those running independent programmes (1.29) on a bug bounty platform ($p = 0.0186$).

7.2. Open-ended results

Within the open-ended answers provided by respondents in the survey and interviews, there are several relevant themes identified. These themes are further separated into concepts pertaining to the operation of CVD programmes in general, and concepts involving the use of third-party platforms. The themes are: general recommendations and advice (Section 7.2.1), overarching benefits (Section 7.2.2), general economic considerations (Section 7.2.3), and the use of platforms (Section 7.2.4).

7.2.1. General recommendations

Respondents provided recommendation and advice that may benefit those looking to launch a CVD programme. Within the theme of general recommendations and advice, four sub-themes emerged: internal recommendations and considerations, setting the programme scope, defining programme policy, and usage of rewards.

As outlined in the BSIMM framework, CVD programmes represent a security activity that may be best suited to organisations with a mature security posture — and, as highlighted in Section 5.2, one that can at times place significant burden on security teams. Respondents highlighted the need for maturity prior to launch: "nobody should start a CVD program before they've invested deeply internally on security" [P20]. Furthermore, small companies may face additional difficulties due to the limited number of personnel available to operate a programme: "I think you probably need a security team of a certain size so that you have resource. If you are a tiny company you may struggle to operate, but then at the same time you may not get many reports because you only have a small web estate." [P3]. Naturally, once a programme is launched, organisations may need to actively refine internal processes and external communications: "there's a bunch of maturity to achieve after establishing some sort of point of contact" [P1].

As mentioned in Section 5.3, it is important to have a proactive view of security throughout an organisation: "in order for a CVD program to be effective, there needs to be high-level institutional support for actually resolving security issues. Otherwise, it will be very difficult to assign and track the work." [P21]

For an organisation about to embark on the journey of establishing a programme, it is recommended that responsibilities are clearly delegated within the CVD team: "I think you need to understand who is going to triage and what, and how they identify who is responsible for fixing on the things you've put in scope" [P3]. This may help to alleviate the issues surrounding project and programme management, as highlighted in Section 5.2. Several respondents recommended that the CVD team should be led by an individual with a background in security and preferably with prior CVD experience, either within an organisation or as a hacker: "if you want to launch a BBP you need a person that knows about the topic because a lot of times this [lack of experience] is bad side of the programme" [P30]. One respondent noted that they "hired a former security researcher to manage this for us" [P7]. A lack of both managerial and CVD specific experience hindered the success of programmes for several respondents.

In the absence of individuals with prior experience, respondents often chose to launch their programmes with limited scope: "if you are unfamiliar with this process, and if you don't have a reasonable scope in your program this can be a massive waste of your time and resources at first" [P19]. It is recommended that the initial scope focus on "high value components (container escape, prod RCE, etc.) and not minor misconfigurations (missing HSTS header)" [P18]. Opening up all assets to hackers while having immature and untested triage, communication, and fixing processes may overwhelm the responsible teams, and infuriate both internal and external stakeholders. In summary, make sure that the scope is "very well defined" [P26], and "start small" [P33].

Further, to define a suitable scope for a CVD programme, operators must also publish the "rules of engagement" (Laszka et al., 2018) — a set of externally-facing clauses explaining the many facets of a programme. Typically, this includes: a boilerplate com-

pany statement, details of any reward structures (e.g. $ payout amounts per CVSS score), assets deemed to be in- and out-of-scope, and Safe Harbor statements. Many respondents opted to survey the content currently used across existing programmes, seeking inspiration as to the important points that should be included: "[we] were looking at people like Google and what they had in their policies and what they were allowing to be submitted, we nicked a bunch of language from other people's disclosure policies it's fair to say" [P3].

Two respondents chose to use open-source templates to help form their initial policies. One highlighted the need for renewed attention in updating the existing templates to reflect up-to-date norms: "I also looked to a few open source initiatives, I think they are probably a bit stale now" [P9]. Others noted that hackers complained about the initial Safe Harbor clauses used in a programme: "they said — we are not happy with your Safe Harbor clauses, we don't think they are safe" [P3].

7.2.2. Overarching benefits

As indicated in Section 7.1, the general perception of CVD programmes was positive amongst the respondents. Several respondents provided further evidence for this in their closing remarks. In particular, it was highlighted that the operation of a CVD programme can allow "access to some of the global top 100 researchers" [P3], some of whom "take their time getting really familiar with your products, and they are not just looking for quick results" [P1], resulting in a "breadth and depth [of vulnerabilities being discovered] which simply wouldn't otherwise be achievable" [P4]. Many emphasised the uniqueness of the issues identified, including "classes of bugs that pen testers don't look for" [P3] and resulting in "paying far less for much better data" [P26].

7.2.3. Economic considerations

In addition to the discussion around setting bounty levels, there are several key economic themes identified that relate to leveraging the work of hackers, including cost effectiveness, influence on the market for vulnerabilities, and issues concerning the hacker 'labour' market.

A prominent theme in the responses was that of cost effectiveness when compared to alternative security investments, such as the expansion of the security team or the usage of external security services. Particularly for smaller businesses in developed countries, CVD programmes can offer security benefits at a cost that is less than the marginal cost of internal expansion (Walshe and Simpson, 2020). One respondent noted that, "[I had] a positive view of the ROI on running one of these programmes, so its not purely a commercial decision but as a small team we have quite limited resources and the necessary the deep expertise so opening it up to the world and encouraging people to support us helped me grow my virtual team as it were" [P9]. For other respondents, their CVD operating costs were less than the cost of external audits and penetration testing services, and in many cases seemingly offered greater returns: "I've heard people say that a bug bounty programme is not a replacement for pen test. I get where they are coming from but the reality in our experience, and I would venture for many companies, a bug bounty programme is better than a pen test." [P26]. Furthermore, some respondents used their CVD programmes as a means to recruit hackers to private, and somewhat informal, penetration testing sessions: "[the programme is a] recruitment mechanism for contracting independent penetration testers and possible full time in house red teams" [P20], allowing the operator to gain the benefits of an individual already familiar with their systems for a fraction of the cost of a professional service.

In addition to discussions surrounding the markets for vulnerabilities, respondents also commented on the existing structure and incentives of the hacker 'labour' market. Reports published by platform owners highlight the geographic distribution of their registered user base (of hackers), demonstrating the global community they are able to leverage (HackerOne, 2019; 2021). Significant numbers of hackers are based in the United States of America, Russia and India, with smaller numbers coming from almost 170 other countries (HackerOne, 2021). Having a large user-base located outside of developed, high-wage, countries brings with it great comparative advantage: hackers in lower-income countries will be willing to work for the prospect of far lower rewards than those in higher-income countries: "$500 isn't that much money ... if it takes somebody 10 hours, maybe 15 hours, it is really good for where they are in the world, a lot of these bug testers are in, it seems, Asia" [P26]. In comparison, with the average wage in the U.S. being $70,000 (U.S. in 2020)[5], CVD rewards may appear less appealing for U.S.-based hackers, but more enticing for U.S. organisations: "working in mostly U.S. and Canadian wages this calculation really works out for us" [P1].

Some respondents spoke negatively of the current labour market structure involving the use of hackers: "bug bounty programs are 'gig economy' and somewhat exploitative especially if they allow focus [to reject reports based] on minor errors" [P18].

7.2.4. Platforms and their services

Respondents offered a plethora of advice regarding the use of platforms and managed programmes, covering: platform advantages and disadvantages, platform pricing, managed programme advantages and disadvantages. As stated previously, hosting a CVD programme on a third-party platform often comes with many built-in services that ease the operation of a programme. Especially when setting up a programme, respondents noted the usefulness of these services: "they provide you some guidelines, like okay according to your industry, this is like the standard, when not the standard, this is the average from other companies that are in your industry and maturity level" [P30]. Furthermore, integrated payment and reward systems shift the burden of legal checks and payment processing away from the operator: "I like that [platform] just does it and its, they do the points and if we do money, great, its easier that way than doing it ourselves. I wouldn't want to go back to doing it ourselves just because it is annoying to deal with that stuff." [P26].

Respondents were most vocal about the prohibitively high costs associated with hosting on a platform and making use of their services: "They are not too cheap, and I'm thinking if you are a small company and you want a bug bounty programme, it is not really viable because it is not cheap. Just the fee for [hosting on] the platform is not cheap." [P30]. For responding in smaller organisations, the sentiment was echoed as it was not always justifiable to pay these expenses: "we also expect our cost of using a third-party tool to be greater than our cost of running the programme (even with the high volume of invalid reports)" [P9].

7.3. Analysis

The overarching consensus amongst respondents was that CVD programmes are beneficial to security, a good use of resources, and a highly-recommended security activity that organisations can adopt to help secure their assets in a cost-effective manner. Using the results of Sections 5 and Section 6, as well as those of this

[5] " ... average wages are obtained by dividing the national-accounts-based total wage bill by the average number of employees in the total economy, which is then multiplied by the ratio of the average usual weekly hours per full-time employee to the average usually weekly hours for all employees" (Organisation for Economic Co-operation and Development (OECD, 2022)

section, we are in a position to make a list of recommendations for those looking to establish a CVD programme.

7.3.1. Pre-deployment recommendations

1. Deploy experienced staff. One of the issues reported by respondents in Section 5.2 and by Al-Banna et al. (2018) is that of inexperienced staff within the CVD or triaging team. Two respondents had a background that involved participating in CVD programmes, and this undoubtedly gave them significant insight into the practice as a whole. It should certainly not be the case that staff have no cybersecurity background (see Section 5.2.3).
2. Restrict initial scope. A common countermeasure to a high volume of low-quality, low-value, reports involves restricting the scope of a programme to only cover those assets deemed to be of higher importance (Section 5.3). However, as recommended by respondents, organisations should start off with a narrow scope to prevent issues from arising after launch (Section 7.2.1).
3. Reward only valuable reports. In conjunction with the previous recommendation, respondents often chose to exclude certain classes of vulnerabilities from reward structures, instead opting to only pay out for higher value reports. This is commonly achieved by excluding low (and sometimes medium) severity vulnerabilities. Although no quantitative advice is provided in order to help set bounty levels, the prevailing comments suggest setting the initial bounty levels to an amount that is sustainable for a given budget. This is clearly an area that would benefit from further quantitative analysis. Furthermore, given the complexities around payout bounties to hackers (Section 7.2.1), it should be decided whether monetary, non-monetary, or no rewards will be offered to hackers.
4. Define internal information pathways. To ease the operation of the programme once launched, it is sensible to have a well-defined system in place for tracking the status of report vulnerabilities (Section 5.2.3), finding the responsible (internal or external) product owners (Section 5.2.3), and communicate patterns of reports (as part of a rudimentary BBLC) to relevant parties (Section 6.3).
5. Seek support from upper management, legal, and product teams. As noted in Section 5.2.3, if internal teams are unwilling to fix vulnerabilities, the benefits to operating a CVD programme are restricted. Further, without authorisation from the legal team, a programme may be unable to operate (Section 7.2.1).

7.3.2. Post-deployment recommendations

1. Modify reward structures to incentivise hackers. An issue faced by respondents was that of low hacker motivation (particularly after the discovery of 'low-hanging fruit' (Section 5.2)), which limits the number of reports that operators receive. Some respondents suggested increasing current payouts as a solution to this issue (Section 7.2.1). The anecdotal evidence used by organisations to increase bounty levels prompts further quantitative analysis into the effect raising bounty levels has on the returns of the programme.
2. Use incoming report information for reactive security. A majority of respondents used the results from their CVD programme to include (shallow) changes within their organisation (Section 6.3). For an organisation not currently utilising their reports, it may be beneficial to use them in a reactive manner to influence the security activities outlined in Section 6.1.
3. Use incoming report information for proactive security. For an organisation with a mature CVD programme, the deeper integration of report information into organisational security strategies may provide further benefit. As detailed in Section 6.2.4, only a few respondents reported using such an approach.

8. Conclusion

Motivated by the question *How effectively are organisations utilising CVD programmes?*, we have considered topics surrounding the use of CVD programmes within organisations via 39 survey responses and 8 semi-structured interviews. A summary of the findings to the research questions presented in Section 2.3 are as follows.

1. Building on the work of Al-Banna et al. (2018), we found that many of the previously identified problems associated with the operation of a CVD programme persist: organisations still struggle with high volumes of poor quality reports that burden the operators and provide little benefit to security. Respondents suggest that alterations to the definitions of in-scope and out-of-scope assets can help to mitigate these issues.
2. Although a majority of organisations report making changes to security activities based upon the content of vulnerabilities reports, there is little evidence suggesting that organisations deeply integrate the results into a form of BBLC. Respondents typically did not institute change throughout their SDLC to the extent of organisations such as Verizon Media.
3. Overall, respondents highly recommended the use of CVD programmes as a cost-effective security activity. Both pre- and post-deployment recommendations were offered to those looking to set up a CVD programme. Designing reward structures to incentivise hackers and avoid rewarding low-value reports, restricting programme scope, and having support from experienced staff and management were all reported to be important considerations.

It is important to recognise the study's limitations. The large non-response rate may limit the generalisability of the findings as only a small proportion of the population could be sampled. In addition, the population sample is biased towards organisations based in the U.S., the U.K., and Canada. The interpretation of the qualitative results, and the answers provided by respondents, may be subject to additional cognitive biases. Confirmation bias, hindsight bias, and perhaps a tendency to portray an organisation's activities in an overly positive light, may impact the conclusions that are drawn from the results.

In conclusion, we hope that what has been presented in this paper is beneficial to those with an interest in the theory and practice of CVD programmes.

Declaration of Competing Interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Thomas Walshe reports financial support was provided by National Cyber Security Centre.

CRediT authorship contribution statement

T. Walshe: Conceptualization, Methodology, Investigation, Data curation, Writing – original draft. **A.C. Simpson:** Writing – review & editing, Supervision, Funding acquisition.

Acknowledgement

The authors would like to thank the reviewers for their insightful and helpful comments. This research was undertaken as part of the Data and Models for Secure Software Engineering project, funded by the UK's National Cyber Security Centre. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

ACM, 2022. Code of ethics and professional conduct. https://www.acm.org/code-of-ethics, Last accessed: February 03, 2022.
Ahmed, A., Lee, B., 2020. Organizational learning on bug bounty platforms. In: Proceedings. 2020 26th Americas Conference on Information Systems, AMCIS 2020. AIS, p. 33.
Al-Banna, M., Benatallah, B., Barukh, M.C., 2016. Software security professionals: Expertise indicators. In: 2016 IEEE 2nd International Conference on Collaboration and Internet Computing (CIC). IEEE, pp. 139–148.
Al-Banna, M., Benatallah, B., Schlagwein, D., Bertino, E., Barukh, M.C., 2018. Friendly hackers to the rescue: How organizations perceive crowdsourced vulnerability discovery. In: Pacific Asia Conference on Information Systems (PACIS), p. 230.
Alomar, N., Wijesekera, P., Qiu, E., Egelman, S., 2020. "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), pp. 319–339.
Alphabet, 2022. Google Play Security Reward Program (GPSRP). https://bughunters.google.com/about/rules/5604090422493184, Last accessed: February 03, 2022.
Ardi, S., Byers, D., Meland, P.H., Tondel, I.A., Shahmehri, N., 2007. How can the developer benefit from security modeling? In: The Second International Conference on Availability, Reliability and Security (ARES'07). IEEE, pp. 1017–1025.
Beckers, K., Heisel, M., Hatebur, D., 2015. Pattern and security Requirements, Engineering-based Establishment of Security Standards. Springer.
Braun, V., Clarke, V., 2006. Using thematic analysis in psychology. Qual Res Psychol 3 (2), 77–101.
Bryman, A., 2016. Social research methods. Oxford University Press.
Buchanan, E.A., Hvizdak, E.E., 2009. Online survey tools: ethical and methodological concerns of human research ethics committees. Journal of Empirical Research on Human Research Ethics 4 (2), 37–48.
Bugcrowd, 2022. Managed bug bounty. https://www.bugcrowd.com/products/bug-bounty/, Last accessed: February 12, 2022.
Bugcrowd, 2022. Public bug bounty program list. https://www.bugcrowd.com/bug-bounty-list/, Last accessed: February 03, 2022.
Catota, F.E., Morgan, M.G., Sicker, D.C., 2019. Cybersecurity education in a developing nation: the Ecuadorian environment. Journal of Cybersecurity 5 (1), tyz001.
Cavusoglu, H., Cavusoglu, H., Raghunathan, S., 2007. Efficiency of vulnerability disclosure mechanisms to disseminate vulnerability knowledge. IEEE Trans. Software Eng. 33 (3), 171–185.
Cobalt, 2022. Customer list. https://cobalt.io/customers, Last accessed: February 03, 2022.
Conklin, W.A., Cline, R.E., Roosa, T., 2014. Re-engineering cybersecurity education in the US: an analysis of the critical factors. In: 2014 47th Hawaii International Conference on System Sciences. IEEE, pp. 2006–2014.
Crumpler, W., Lewis, J.A., 2019. The cybersecurity workforce gap. Center for Strategic and International Studies (CSIS) Washington, DC, USA.
Cuevas, A., Hogan, E., Hibshi, H., Christin, N., 2022. Observations from an online security competition and its implications on crowdsourced security. arXiv preprint arXiv:2204.12601.
Desselle, S.P., 2005. Construction, implementation, and analysis of summated rating attitude scales. Am J Pharm Educ 69 (5).
Disclose.io, 2022. Program database. https://disclose.io/programs/, Last accessed: February 03, 2022.
Ellis, R., Stevens, Y.,. Bounty everything: Hackers and the making of the global bug marketplace. Available at SSRN 4009275, 2022, SSRN. 2022
Emami-Naeini, P., Agarwal, Y., Cranor, L.F., Hibshi, H., 2020. Ask the experts: What should be on an IoT privacy and security label? In: 2020 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 447–464.
Etcovitch, D., van der Merwe, T., 2018. Coming in from the cold: A safe harbor from the CFAA and the DMCA ú1201 for security researchers (June 1, 2018). Berkman Klein Center Research Publication No. 2018-4, SSRN.
Fan, Y., Xia, X., Lo, D., Hassan, A.E., 2018. Chaff from the wheat: characterizing and determining valid bug reports. IEEE Trans. Software Eng. 46 (5), 495–525.
Fink, A., 2003. The survey handbook. SAGE Publications, Inc..
FireBounty, 2022. The right path to coordinated vulnerability disclosure. https://firebounty.com, Last accessed: February 03, 2022.
Follis, L., Fish, A., 2022. State hacking at the edge of code, capitalism and culture. Information, Communication & Society 25 (2), 242–257.
Formosa, P., Wilson, M., Richards, D., 2021. A principlist framework for cybersecurity ethics. Computers & Security 109, 102382.
Furnell, S., Fischer, P., Finch, A., 2017. Can't get the staff? the growing need for cyber-security skills. Computer Fraud & Security 2017 (2), 5–10.
Galesic, M., Bosnjak, M., 2009. Effects of questionnaire length on participation and indicators of response quality in a web survey. Public Opin Q 73 (2), 349–360.
HackerOne, 2019. The 2019 hacker report. https://www.hackerone.com/sites/default/files/2019-02/the-2019-hacker-report_3.pdf, Last accessed: September 23, 2019.
HackerOne, 2021. The 2020 hacker report. https://www.hackerone.com/resources/reporting/the-2020-hacker-report, Last accessed: February 10, 2021.
HackerOne, 2022. Product offerings. https://docs.hackerone.com/programs/product-offerings.html, Last accessed: February 12, 2022.
HackerOne, 2022. Program directory. https://hackerone.com/directory/programs, Last accessed: February 03, 2022.
Hafiz, M., Fang, M., 2016. Game of detections: how are security vulnerabilities discovered in the wild? Empirical Software Engineering 21 (5), 1920–1959.
Hata, H., Guo, M., Babar, M.A., 2017. Understanding the heterogeneity of contributors in bug bounty programs. In: Proceedings of the 11th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement. IEEE Press, pp. 223–228.
Hern, A., 2019. The Guardian: apple to pay hackers more than $1m to find security flaws. https://www.theguardian.com/technology/2019/aug/12/apple-hackers-black-hat-conference, Last accessed: September 10, 2019.
Hilton, M., Nelson, N., Tunnell, T., Marinov, D., Dig, D., 2017. Trade-offs in continuous integration: assurance, security, and flexibility. In: Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, pp. 197–207.
Householder, A.D., Wassermann, G., Manion, A., King, C., 2017. The CERT guide to Coordinated Vulnerability Disclosure. Technical Report. Carnegie-Mellon University.
iNTiGRiTi, 2022. Public bug bounty programs. https://www.intigriti.com/programs, Last accessed: February 03, 2022.
Jones, R.L., Rastogi, A., 2004. Secure coding: building security into the software development life cycle. Information Systems Security 13 (5), 29–39.
Joshi, D., 2022. Awesome bug bounty. https://github.com/djadmin/awesome-bug-bounty, Last accessed: February 03, 2022.
Kapoor, A., Penton, A., Pierpont, H., 2022. Eliciting course feedback through a bug bounty program. In: Proceedings of the 27th ACM Conference onInnovation and Technology in Computer Science Education Vol. 2, pp. 595–596.
Kerlinger, F.N., 1966. Foundations of behavioral research. Holt, Rinehart and Winston: New York.
Krosnick, J.A., 2018. Questionnaire design. In: The Palgrave handbook of survey research. Springer, pp. 439–455.
Laszka, A., Zhao, M., Grossklags, J., 2016. Banishing misaligned incentives for validating reports in bug-bounty platforms. In: European Symposium on Research in Computer Security. Springer, pp. 161–178.
Laszka, A., Zhao, M., Malbari, A., Grossklags, J., 2018. The rules of engagement for bug bounty programs. In: International Conference on Financial Cryptography and Data Security, pp. 138–159.
Li, Y., Zhao, L., 2022. Collaborating with bounty hunters: how to encourage white hat hackers' participation in vulnerability crowdsourcing programs through formal and relational governance. Information & Management 59 (4), 103648.
Lietz, P., 2010. Research into questionnaire design: a summary of the literature. International Journal of Market Research 52 (2), 249–272.
Malladi, S.S., Subramanian, H.C., 2019. Bug bounty programs for cybersecurity: practices, issues, and recommendations. IEEE Software 37 (1), 31–39.
Manfreda, K.L., Bosnjak, M., Berzelak, J., Haas, I., Vehovar, V., 2008. Web surveys versus other survey modes: ameta-analysis comparing response rates. International Journal of Market Research 50 (1), 79–104.
Manjikian, M., 2017. Cybersecurity ethics: an introduction. Routledge.
Media, V., 2022. Brighttalk webinar: Breaking barriers: Introducing the bug bounty lifecycle. https://core.brighttalk.com/webcast/13109/458461, Last accessed: January 27, 2022.
O'Hare, J., Shepherd, L.A., 2022. Developing a gamified peer-reviewed bug bounty programme. In: International Conference on Human-Computer Interaction. Springer, pp. 514–522.
Poteat, T., Li, F., 2021. Who you gonna call? an empirical evaluation of website security.txt deployment. In: Proceedings of the 21st ACM Internet Measurement Conference. ACM, pp. 526–532.
Pressman, R.S., 2005. Software Engineering: A Practitioner's Approach. Palgrave Macmillan.
Sanchez, M.E., 1992. Effects of questionnaire design on the quality of survey data. Public Opin Q 56 (2), 206–217.
Sappleton, N., Lourenço, F., 2016. Email subject lines and response rates to invitations to participate in a web survey and a face-to-face interview: the sound of silence. Int J Soc Res Methodol 19 (5), 611–622.
Schuman, H., Presser, S., 1979. The open and closed question. Am Sociol Rev 692–712.
Silic, M., Lowry, P.B., 2021. Breaking bad in cyberspace: understanding why and how black hat hackers manage their nerves to commit their virtual crimes. Information Systems Frontiers 23 (2), 329–341.
Silomon, J., Hansel, M., Schwartz, F., 2022. Bug bounties: between new regulations and geopolitical dynamics. In: International Conference on Cyber Warfare and Security, Vol. 17, pp. 298–305.
Smith, J., Theisen, C., Barik, T., 2020. A case study of software security red teams at Microsoft. In: 2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC). IEEE, pp. 1–10.
Sullivan, G.M., Artino Jr, A.R., 2013. Analyzing and interpreting data from likert-type scales. J Grad Med Educ 5 (4), 541–542.

Synopsys, 2022. BSIMM12 2021 Foundations report. https://www.bsimm.com/content/dam/bsimm/reports/bsimm12-foundations.pdf, Last accessed: March 4, 2022.
Synopsys, 2022. BSIMM12: Building Security in Maturity Model version 12. https://www.bsimm.com/download/, Last accessed: March 4, 2022.
Synopsys, 2022. BSIMM6: Building Security in Maturity Model version 6. https://www.inf.ed.ac.uk/teaching/courses/sp/2015/lecs/BSIMM6.pdf, Last accessed: March 4, 2022.
Synopsys, 2019. BSIMM9: Building Security in Maturity Model version 9. https://www.bsimm.com/download/, Last accessed: December 18, 2019.
Takanen, A., Demott, J.D., Miller, C., Kettunen, A., 2018. Fuzzing for software security testing and quality assurance. Artech House.
Tanczer, L.M., 2020. 50 Shades of hacking: how IT and cybersecurity industry actors perceive good, bad, and former hackers. Contemporary Security Policy 41 (1), 108–128.
International Organization for Standardization, 2022. ISO/IEC 29147:2018: Information technology, security techniques, vulnerability disclosure. https://www.iso.org/standard/72311.html, Last accessed: March 4, 2022.
Organisation for Economic Co-operation and Development (OECD, 2022. Average wages. https://data.oecd.org/earnwage/average-wages.htm, Last accessed: March 4, 2022.
U.S. Department of Homeland Security, 2022. Binding operational directive 20-01. https://cyber.dhs.gov/assets/report/bod-20-01.pdf, Last accessed: March 4, 2022.
Vandervelden, S., Chowdhury, M.M., Latif, S., 2021. Managing the cyber world: Hacker edition. In: 2021 International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME). IEEE, pp. 1–6.
Votipka, D., Stevens, R., Redmiles, E., Hu, J., Mazurek, M., 2018. Hackers vs. testers: a comparison of software vulnerability discovery processes. In: 2018 IEEE Symposium on Security and Privacy (SP). IEEE, pp. 374–391.
Wachs, J., 2022. Making markets for information security: the role of online platforms in bug bounty programs. arXiv preprint arXiv:2204.06905.
Walshe, T., Simpson, A., 2020. An empirical study of bug bounty programs. In: 2020 IEEE 2nd International Workshop on Intelligent Bug Fixing (IBF). IEEE, pp. 35–44.
Walshe, T., Simpson, A., 2022. A longitudinal study of hacker behaviour. In: The 37th ACM/SIGAPP Symposium on Applied Computing (SAC '22). ACM, pp. 1465–1474.
Witschey, J., Zielinska, O., Welk, A., Murphy-Hill, E., Mayhorn, C., Zimmermann, T., 2015. Quantifying developers' adoption of security tools. In: Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, pp. 260–271.
Xia, X., Lo, D., Ding, Y., Al-Kofahi, J.M., Nguyen, T.N., Wang, X., 2016. Improving automated bug triaging with specialized topic model. IEEE Trans. Software Eng. 43 (3), 272–297.
YogOsha, 2022. Our clients. https://yogosha.com/yogosha-clients/, Last accessed: February 03, 2022.
Zhao, M., Grossklags, J., Chen, K., 2014. An exploratory study of white hat behaviors in a web vulnerability disclosure program. In: Proceedings of the 2014 ACM Workshop on Security Information Workers. ACM, pp. 51–58.
Zhao, M., Grossklags, J., Liu, P., 2015. An empirical study of web vulnerability discovery ecosystems. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. ACM, pp. 1105–1117.
Zou, W., Lo, D., Chen, Z., Xia, X., Feng, Y., Xu, B., 2018. How practitioners perceive automated bug report management techniques. IEEE Trans. Software Eng. 46 (8), 836–862.
Zrahia, A., Gandal, N., Markovich, S., Riordan, M.H., 2022. The simple economics of an external shock on a crowdsourced 'bug bounty platform'. SSRN. Available at SSRN 4154516, 2022

Thomas Walshe is a D.Phil. candidate at the University of Oxford. Under the supervision of Prof. Andrew Simpson, he is researching how CVD programmes can support data-driven software development lifecycles. This topic is being pursued as part of the Data and models for secure software engineering project, which is funded by the U.K.'s National Cyber Security Centre.

Andrew Simpson is an Associate Professor in Software Engineering at the University of Oxford. He is the Principal Investigator of the aforementioned Data and models for secure software engineering project.