Accepted: 27 October 2025
Revised: 22 September 2025
Received: 20 May 2025

RESEARCH-ARTICLE

Vulnerability Disclosure or Notification? Best Practices for Reaching Stakeholders at Scale

TINGHAN CHEN, University of Twente, Enschede, Overijssel, Netherlands

JEROEN VAN VOS, University of Twente, Enschede, Overijssel, Netherlands

Digital Threats: Research and Practice
https://doi.org/10.1145/3777485
EISSN: 2576-5337

Vulnerability Disclosure or Notification? Best Practices for Reaching Stakeholders at Scale

TING-HAN CHEN[*] and JEROEN VAN DER HAM-DE VOS[*], University of Twente, The Netherlands

Vulnerability disclosure is the practice of a finder disclosing a newly found vulnerability to vendors. It has received best practices to ensure communication between stakeholders. However, the practice of a finder or vendor notifying end-users about vulnerable systems and mitigation plans has not received the same attention and guidelines for performing it at scale. We identify the practice as vulnerability notification, which shares similarities with disclosure but presents other challenges and requires different approaches.

In vulnerability notification, a finder targets known vulnerabilities or misconfigurations using active scans or datasets to determine how many systems or services remain vulnerable. The scale and complexity of vulnerability notification to end-users are often significantly greater than those of multi-party disclosure to vendors. These place an increasing burden on finders to inform stakeholders on time, especially for academic security researchers, ethical hackers, and practitioners.

Based on our experience with notifications and academic publications documenting disclosure and notification operations, we conduct a meta-review of how researchers have adopted best practices, pursued different strategies, and reflected on their operations over the years. Drawing on the meta-review and suggestions from security communities, we propose new best practices for finders to perform vulnerability disclosure, particularly vulnerability notification at scale.

CCS Concepts: • **Security and privacy → Vulnerability management**.

Additional Key Words and Phrases: Vulnerability disclosure, Vulnerability notification, Best practice

1 INTRODUCTION

Coordinated Vulnerability Disclosure (CVD), formerly known as responsible disclosure, is the accepted best practice in the security community for handling discovered vulnerabilities. When a new vulnerability is found, the finder conducts a risk and impact assessment and then initiates conversations with stakeholders to mitigate vulnerable systems or services in time. However, due to shifts in the network landscape, the scale of vulnerable systems and the number of stakeholders involved have changed drastically from the past [9, 41]. The amount and variety of vulnerabilities have increased [30], and so have the affected parties [3, 25]. This presents a new challenge for finders, particularly academic security researchers and practitioners: figuring out a scalable method to identify reliable contact information and ensure message delivery to stakeholders.

Over the years, support mechanisms, such as vulnerability disclosure policies [32], the VINCE vulnerability platform [5], and bug bounty programs [46], have come into place to address the challenge of vulnerability disclosure. Finders and vendors can adopt CVD to perform disclosure with the support organisations such as Computer Security Incident Response Teams (CSIRTs) [14] and Product Security Incident Response Teams (PSIRTs) [16]. However, not every finder or vendor has the same capacity and experience to conduct disclosure at

[*] Both authors contributed equally to this research.

Authors' Contact Information: Ting-Han Chen, t.h.chen@utwente.nl; Jeroen van der Ham-de Vos, j.vanderham@utwente.nl, University of Twente, Enschede, Overijssel, The Netherlands.

ACM 2576-5337/2025/11-ART
https://doi.org/10.1145/3777485

scale [47]. Moreover, finders from different communities, such as academic security researchers [3], individual ethical hackers, practitioners, and bug bounty hunters, may have different security interests and constraints in selecting their approaches to reach stakeholders.

Despite challenges in vulnerability disclosure, notifications to stakeholders about known vulnerabilities and security issues in existing systems have gained attention over the years [41]. We identify the practice as vulnerability notification, which informs end-users, such as product owners, hosting providers, domain owners, network operators, and incident responders. These stakeholders differ from vendors, who are traditionally stakeholders in CVD. Whether the vulnerability is known to stakeholders and the public is the main difference between vulnerability disclosure and vulnerability notification. In *vulnerability disclosure*, a finder discovers a new vulnerability and plans to disclose it to the vendor. Meanwhile, in *vulnerability notification*, a finder identifies a known vulnerability on existing machines and notifies end-users. This means that vulnerability notification is often conducted after the vulnerability disclosure. In particular, after a newly found vulnerability has been disclosed and mitigated with stakeholders or even made public, it is not guaranteed that all vulnerable systems are treated on a timely basis [9]. Hence, vulnerability notifications are intended to improve Internet safety.

Vulnerability notification is often initiated after a finder performs network scanning [40] or vulnerability analysis of existing datasets [6] to identify vulnerable systems with targeted vulnerabilities. A finder should retrieve the stakeholders' contact information and select a communication channel to inform the affected parties. However, similar to the challenge in vulnerability disclosure, notification to stakeholders also suffers from the high number and complexity of stakeholders [35, 50]. The contact retrieval and notification at scale pose an even tougher challenge for finders [26, 37], since stakeholders in vulnerability notification may involve a more diverse and complex set of parties [6, 50].

1.1 Research Questions

How to disclose a vulnerability to a vendor is well considered, with CVD as the best practice and increased support mechanisms suggested for finders, governments, and vendors [13]. However, it is still unclear how well the best practices and mechanisms are known in the academic security research community. Furthermore, the lack of best practices for vulnerability notification has posed a growing challenge not only for academic security researchers but also for other stakeholders. This motivates us to compile the research questions as follows:

(1) What distinguishes vulnerability disclosure and notification, especially in communication to stakeholders?
(2) How did researchers adopt best practices to carry out vulnerability disclosure and notification at scale?
(3) What insights can we gain from the researchers' experiences of large-scale disclosure and notification?
(4) What are the best practices for researchers and finders to perform vulnerability disclosure and notification?

In the following sections, we aim to answer each question with insights we collect from the literature, security communities, support organisations, and our own notification experience. In Section 2, we distinguish between vulnerability disclosure and notification, and explain how a large-scale scenario can affect the two distinct practices in operations. In Section 3, we explain how we select the literature and conduct a meta-review using our proposed stage model to extract experiences from the academic security research community. In Section 4, we compile the insights we gather from the academic experience in each stage of vulnerability disclosure and notification. Finally, in Section 5, we propose our best practices for the finders to address the current limitations of large-scale vulnerability disclosure and notification. Best practices can also help other stakeholders improve their disclosure and notification processes.

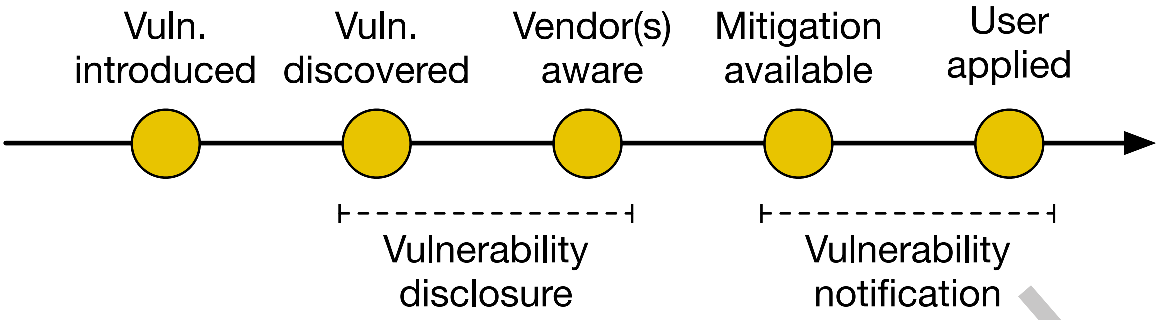

Fig. 1. Vulnerability Disclosure and Notification

2 BACKGROUND AND RELATED WORK

2.1 Vulnerability Disclosure and Notification

Vulnerability disclosure, especially coordinated vulnerability disclosure, has been an accepted best practice in security research for years. There have been numerous publications and even standards published on this process [13]. However, notifying users of vulnerable systems has usually not been seen as a separate practice and can come with its own challenges. The difference between the two processes is illustrated in Figure 1.

In *vulnerability disclosure*, a finder discovers a new vulnerability and aims to share this with the vendor. The role of the finder in the following content includes both the responsibilities of identifying a vulnerability and informing the responsible parties [15]. Once the disclosure process has started, the vulnerability is usually revealed for the first time. The finder and vendor(s) discuss the finding with or without a coordinator's support. Subsequently, the vendor works on a mitigation. Vulnerability disclosure requires trusted channels and detailed message information to minimise disclosure leaks to unintended recipients. The finder may need to expect additional steps to communicate with the stakeholder about whether and how to disclose the vulnerabilities to the public, based on the disclosure policy and the legal agreement on both sides. After some time, the vulnerability is made public, usually with a possible mitigation available.

In *vulnerability notification*, the vulnerability information is already available to the stakeholders before a notification starts. This means the vulnerability has already been documented and revealed to the stakeholders or the public. The weakness is likely documented with CVE numbers or discussed within specific communities. Potential malicious parties, such as criminals, may also be aware of the vulnerabilities and attempt to exploit them. The receiving users may or may not be aware of the vulnerability's existence before seeing the notification message. Once they are aware of the issues, they may be able to find the vulnerability information with resources that are not limited to the finder. This can lead to different user behaviours in response to the finder. The affected user may also have a different risk assessment based on the severity of the vulnerabilities. The aim of the notification process is to improve general security by removing vulnerable systems.

We note that there are many overlapping challenges in vulnerability disclosure and notification, including formulating the initial message, choosing the communication channel, and handling responses. However, due to the nature of the different stakeholders involved and especially the differences in scale, the challenges are fundamentally different.

2.2 Evaluations of Disclosure and Notification

Vulnerability disclosure has received best practices over the years through contributions from security researchers, vendors, and support organisations. There have been several publications that describe the practices in detail and examine their effectiveness. Householder and Spring [20] proposed a model to assess the coordinator role in CVD, providing insight into vendor incident handling and the vulnerability lifecycle. Walshe and Simpson [47] examined the effectiveness of how vendors with CVD programs and outsourced vulnerability platforms receive and process vulnerability disclosure. They revealed that disclosure program operators receive a large number of vulnerability reports, which creates a burden for vulnerability prioritisation. Nakajima et al. [30] examined vulnerability management among IoT vendors across two countries and identified disclosure pitfalls to avoid.

Vulnerability notification has received growing attention over the last decade, especially regarding its efficiency. The Dutch Institute for Vulnerability Disclosure (DIVD) [10] proposed the notification guideline [41] based on its framework and the Communication-Human Information Processing (C-HIP) model [48], drawing on its notification operational experience. Their contribution focuses on scalable notification strategies for end-users, such as an incident responder and an abuse specialist. Additionally, the challenges they documented with end-users resemble the capacity and awareness issues faced by program operators discussed by Walshe and Simpson in the context of vulnerability disclosure [47].

Evaluations of vulnerability disclosure and notification at scale still require attention, despite contributions that have sought to address the challenges from different angles over the years. An extensive assessment and best practices for performing disclosure, especially notification at scale, are needed. There has been research focusing on finders trying to reach out to other stakeholders effectively and documenting the complete process of their disclosure and notification operations. The work will be discussed in the following sections with our assessment and comparison with best practices suggested by the communities and support organisations.

2.3 Large-Scale Disclosure and Notification

Large-scale disclosure operations have been reported as an increasing challenge by studies in the last decade [20, 30, 47]. Unlike the one-to-one or one-to-multiple disclosure to vendors, the complexity and number of stakeholders to inform have increased to a level beyond the effort individuals and small teams can make.

Many of the challenges associated with multi-party disclosure processes have been identified by CERT/CC [21] and the FIRST Special Interest Group on Vulnerability Coordination [17], resulting in a best-practice document, "Guidelines and practices for Multi-Party Vulnerability Coordination and Disclosure" [15]. These best practices address the challenges of disclosing newly discovered vulnerabilities and, through various scenarios, describe the challenges and impacts for different stakeholders, including finders, vendors, defenders, and users. The VINCE platform [5] has been developed by CERT/CC and other contributors to support multi-party disclosures and prevent many of the possible mistakes identified in the best practice guide.

Moreover, many countries have introduced laws, regulations [11], and vulnerability disclosure policies [8, 31] to stimulate stakeholders, with active support from national CSIRTs [14] and PSIRTs [16], professional organisations [13, 15], and others, all supporting the practice of vulnerability disclosure. Likewise, research institutes, universities, and companies, such as Google and Facebook, have adopted outgoing disclosure policies to protect the finders and make their intentions clear to receiving parties. [32, 41].

A vulnerability disclosure to multiple parties is considered to be large-scale when there are five or more vendors involved, and can become very complicated and stressful [3, 29, 42]. However, vulnerability *notification* processes can quickly scale up, with affected parties numbering in the hundreds or even thousands. Identifying the contact information and notifying the responsible parties behind the vulnerable systems can be overwhelming for finders [26, 35, 36, 50], which can eventually limit the development of vulnerability notification.

There have been contributions to improve the efficiency of the notification mechanism. Emails have been empirically confirmed as the method for reaching large numbers of affected parties, despite the drawbacks of low delivery rates and inaccurate contact information [36, 37, 41, 50]. Nonetheless, the current best practices for vulnerability disclosure, especially vulnerability notification, are still struggling to catch up with the growing nature and importance of digital infrastructure.

With best practices and prior experiences as a foundation for large-scale disclosure, the recent effort still suffers from the slow adoption of disclosure policy and incident handling among stakeholders [6, 42]. Although bug bounty programs, vulnerability disclosure policies, and regulations are adopted as best practices by the industry and countries, not every finder from different communities, such as academic security researcher, practitioner, ethical hacker, and bounty hunter, shares the same capacity and security interest to ensure the message delivery and conversations of the disclosure and notification at scale with the current best practices and guidelines [3, 29, 46]. For instance, academic security researchers may face time constraints in publishing their findings, while practitioners in a small team may have limited capacity and experience to prioritise stakeholders when large-scale scenarios are considered. The current best practices do not necessarily cover the disclosure or notification guidelines for such finders.

There are increasingly large-scale internet-wide vulnerability notification cases beyond the support of local or national support organisations, leading finders, such as academic security researchers, to carry out notifications by setting up their own messaging infrastructure [6, 35]. It was reported that contact identification prioritisation remains challenging, and the notification operation is still not widely accepted among stakeholders [6, 29], even though disclosure policies or vulnerability report programs were in place. This indicates that large-scale notifications still necessitate new proposals to enhance the process for both notifying and receiving parties.

So far, we have answered our first research question in Section 1.1. With a clear distinction between vulnerability disclosure and notification at scale, we aim to fill the gap by conducting a meta-review of academic security researchers' experiences, with suggestions from security communities, support organisations, and our experience.

3 LITERATURE SELECTION

To understand how the academic security research community has approached vulnerability disclosure and notification at scale over time, we collected publications with extensive documentation of large-scale vulnerability disclosure or notification operations over the last decade. We initially used literature search engines to identify publications, then selected those that included disclosure procedures in stages and documented pre- and post-disclosure or notification operations. We did not exhaustively locate all available disclosure or notification work. Instead, we selected work that can be representative of large-scale disclosure and notification implementation across different disclosure and notification scenarios from the perspective of academic security researchers.

The initial search patterns we used were the combination of "vulnerability", "disclosure", "notification", "report", "large-scale", and "network" to look for the publications through major academic security conferences, literature databases, and search engines such as ACM Digital Library, IEEE Xplore, Scopus, and Google Scholar with all fields or metadata search filter enabled. We skimmed the titles, abstracts, and partial content of search results that fit the large-scale disclosure or notification scenarios discussed in Section 2.3. The initial result was broad and diverse. After a few iterations, we narrowed down our search keywords to "vulnerability disclosure" and "vulnerability notification", with synonyms such as "vulnerability alert" and "vulnerability warning". It is worth noting that we also examine publications that don't necessarily include these keywords but do have related content with manual effort. This gave us a short list of approximately 60 results that were relevant to vulnerability disclosure and notification.

To finalise the short list of publications, we examined the literature candidates with details such as assessment before an operation, selection of communication channels and messaging infrastructure, review after the operations, and contributions to best practices. We assigned different categories based on communication channels and involved stakeholders to identify areas of difference and choose work that has a measurable impact on the stakeholders. Furthermore, to understand the development of disclosure and notification best practices over the years, we aimed for publications that can be used to compare with each other, such as adopting, changing or improving methods based on other selected work. This means we looked into the reference list and related work sections of the publications to seek correlations and influences. Besides, we selected work that had an impact on the security communities to illustrate the CVD as a best practice and to stimulate discussion on how we can improve the existing guidelines. Finally, we separated the publications into vulnerability disclosure and notification operations as mentioned in Section 2.1.

As a result, we selected 15 distinct publications that well represent large-scale vulnerability disclosure and notification using several approaches in different scenarios from the last decade. We separate the work into two tables in chronological order. Table 1 presents the vulnerability disclosure operations, whereas Table 2 presents the vulnerability notification operations. To comprehend the selected work, we develop assessment stages to look into each disclosure or notification operation in the following subsection.

3.1 Stages of Operation

In the literature selection, we cross-reviewed each selected work in detail and aimed to figure out the common procedure for performing vulnerability disclosure and notification. Eventually, we compiled five stages to represent the procedure implemented across the literature. We then used the stages to extract key points, numbers, and remarks from the selected work to better understand efforts, considerations and reviews made in the operations. This results in the following stages to assess the selected publications:

Pre-Assessment – Before a vulnerability disclosure or notification can start, a finder will assess the risk and impact of the discovered vulnerability. We identify the type and number of stakeholders involved and the vulnerabilities, and then extract the impact scale of the vulnerability to understand the preparation required before an operation.

Communication Channel – After deciding on the stakeholders to inform, the proper communication channels should be selected to deliver the messages. We list the single to multiple communication channels used in each operation to inform the affected parties.

Messaging Infrastructure – For the message to be delivered with the selected channel, the right messaging infrastructure should be used to ensure message delivery. We distinguish the messaging infrastructure used by finders to deliver the message and the infrastructure used by the stakeholders to receive and forward the message.

Disclosure Policy and Message – The wording of the message should be tuned to the stakeholders to ensure comprehension and follow the needs of the recipients. We reflect on how the finders composed their message and how they handled conversations with other stakeholders. In addition, we also check on disclosure policies used by the finders and other stakeholders, if presented.

Post-Review – An operation can be reviewed by tracking the remediation rate, feedback from stakeholders, and experiences in each stage. A finder can also reflect on the operation and contribute to best practices. We extract the notified hosts and domains, remediation rates, stakeholder interactions, and contributions to best practices for large-scale disclosure and notification operations.

To present an overview of the assessment for each publication, we arrange three tables based on the five stages outlined above. Tables 1 and 2 both include the publication title, pre-assessment, and post-review to indicate the efforts and outcomes for each work, with different scenarios considered. Table 3 shows the single or multiple communication channels and messaging infrastructure implemented in each disclosure or notification operation.

Title	Year	Pre-Assessment	Post-Review
Key Reinstallation Attacks: Forcing Nonce Reuse in WPA2 [44]	2017	To Wi-Fi related vendors, Key reinstallation attacks, At least 6 vendors	Increased to 84 vendors (VINCE), Measurement & notification delegated to CERT/CC & vendors
Talking Trojan: Analysing an Industry-Wide Disclosure [3]	2022	To software vendors, Trojan Source, 13 vendors	Increased to 19 vendors (VINCE), Measurement delegated to vendors, GitHub, and Rust team
rpkiller: Threat Analysis of the BGP Resource Public Key Infrastructure [42]	2023	To software vendors, RPKI implementations, 8+ software vendors	5 vendors released fixes, 2 vendors didn't release the fix, 1 vendor stopped support
Vulnerability Disclosure Considered Stressful [29]	2023	To vendors, network operators TsuNAME,DNS resolver&clients, 5+ vendors, communities	Measurement not specified, 1 DNS community meeting, 3 security events
Haunted by Legacy: Discovering and Exploiting Vulnerable Tunnelling Hosts [2]	2025	To vendors, domain owners, Tunnelling protocol & EDoS, 3,527,565 IPv4, 735,628 IPv6 hosts	14 domains directly notified, Measurement & notification delegated to CSIRTs & Shadowserver

Table 1. Vulnerability Disclosure Operations

The five stages and the three tables serve as our response to the second research question in Section 1.1 and as the foundation for understanding the selected publications and extracting insights in the next section.

4 LESSONS LEARNED

This section presents the insights we have gained from the selected work on large-scale vulnerability disclosure and notification. The main focus is from the finder's perspective and academic research experience. We also integrate the experiences we learn from security communities, support organisations, existing guidelines and standards to present the essence of the disclosure or notification procedure.

With Tables 1, 2, and 3 presented as an overview of disclosure and notification operations, we discuss the development of best practices from the experiences of the selected work in Section 4.1. Then, we compare how the finders adopted the past best practices and what they learned from their operations in Section 4.2.

4.1 The Development of Best Practices

Academic security researchers as finders have been seeking optimal methods and strategies based on the disclosure and notification best practices presented in their time, as shown in Tables 1 and 2. We have observed changes in their communication channels, messaging infrastructure, contact retrieval, message formulation, and disclosure policies, all of which are influenced by the nature of the vulnerabilities found, notification scalability, and community recommendations.

The essential part of vulnerability disclosure and notification is ensuring messages reach the responsible parties and prompt them to action. As presented in Table 3, the finders used different communication channels and messaging infrastructure to ensure the delivery of the disclosure messages, either on their own or with support organisations, such as national CSIRTs and The Shadowserver Foundation. To maximise notification coverage, all work adopted multiple communication channels, and most delegated notifications to support organisations.

The Shadowserver Foundation is a non-profit organisation founded in 2004, driven by the vision of a secure, threat-free internet [34]. They do this by scanning the internet for known vulnerabilities, performing notifications,

Title	Year	Pre-Assessment	Post-Review
The Matter of Heartbleed [9]	2014	To network operators, TLS Heartbeat Extension, 588,686 vulnerable hosts	212,805 hosts notified, 4,648 emails (WHOIS), 57% mitigated
You've Got Vulnerability: Exploring Effective Vulnerability Notifications [25]	2016	To network operators, 45,770 ICS, 83,846 DDoS Ampl., 180,611 IPv6 Firewall hosts	79.7% ICS,92.4% Ampl,99.8% IPv6 notified, 9,918 emails (WHOIS), Up to 18% mitigated
Hey, You Have a Problem: On the Feasibility of Large-Scale Web Vulnerability Notification [37]	2016	To website owners WordPress & Client-Side XSS, 44,790 vulnerable domains	35,832 domains notified, 17,819 emails (alias, WHOIS), 25.8% WP, 12.6% CXSS mitigated
Make Notifications Great Again: Learning How to Notify in the Age of Large-Scale Vulnerability Scanning [50]	2017	To nameserver/network operators & domain owners, DNS Dynamic Updates, 21,506 vulnerable domains	4,149 nameserver IPs notified, 5051 emails (WHOIS), Up to 20% mitigated
Didn't You Hear Me? - Towards More Successful Web Vulnerability Notifications [36]	2018	To website owners, WordPress & Git, 24,000+ vulnerable domains	20,602 domains notified, 103,819 emails (alias, WHOIS), 17% WP, 24% Git mitigated
Tell Me You Fixed It: Evaluating Vulnerability Notifications via Quarantine Networks [49]	2019	To ISP retail customers, DNS resolvers & mDNS services, 1688 retail customers	688 customers notified, 86.6% walled garden only, 75.1% email only mitigated
User Compliance and Remediation Success after IoT Malware Notifications [33]	2021	To ISP retail customers, IoT Malware (Mirai Family), 177 retail customers	From 50 responded participants, 95% walled garden, 82% email only started mitigation
Effective Notification Campaigns on the Web: A Matter of Trust, Framing, and Support [26]	2021	To website owners, IP Anonymisation, 7979 non-compliant sites	4754 sites notified, 2660 letters, 1,337 emails (Manual), 76.3% letter, 33.9% email mitigated
Uncovering the Role of Support Infrastructure in Clickbait PDF Campaigns [35]	2024	To website owners, Clickbait PDF, 177,835 vulnerable hosts	8,842 domains notified, 1,522 emails (alias, WHOIS), 29.567% mitigated
Are You Sure You Want To Do Coordinated Vulnerability Disclosure? [6]	2024	To IoT backend operators, MQTT & Misconfigurations, 15,820 vulnerable backends	15046 backends notified, 2,132 emails (WHOIS), 2.25% mitigated

Table 2. Vulnerability Notification Operations

managing sinkholes for malware command-and-control infrastructure, and operating honeypots and honeyclients to monitor threat developments. Over the years, the team at The Shadowserver Foundation has built an extensive, trusted network with national CSIRTs, ISPs, and others to ensure they can reach as many stakeholders as possible. Anyone can sign up to receive threat reports for their domain or IP address, provided they can prove ownership. More importantly, The Shadowserver Foundation can support researchers in performing notifications to stakeholders, as described by Beitis and Vanhoef [2, 38].

Earlier vulnerability notification operations in [9, 25] focused on email and coordinator-delegation notifications; others in [36, 37, 50] extensively implemented diverse communication channels, empirically compared their

Channel & Infrastructure		Work
Email	Individual Account	[2, 29, 35–37, 42, 44]
	Dedicated Account	[6, 25, 26, 33, 49, 50]
	Not Specified	[3, 9]
Coordinator Delegation		[2, 3, 25, 37, 42, 44]
Website		[2, 3, 6, 25, 44, 50]
Survey		[9, 25, 36, 50]
Phone		[26, 33, 37]
ISP Intervention		[33, 49]
Community & Meeting		[29, 42]
Post		[26, 37]
Social Media		[37]
The Shadowserver Foundation		[2]

Table 3. Channel and Messaging Infrastructure

effectiveness, and provided valuable suggestions for best practices in large-scale notifications. These operations served as the foundation for the large-scale notification best practices in the academic security community.

The recent notification operations in [6, 35] adopted best practices from the prior work. The researchers narrowed down the channel selection and dived into the experiences with email notifications. On the contrary, Maass et al. [26] decided to explore alternatives to best practices that involve costly manual effort via email, posts, and phone calls. All their results reflected the prior best practices and helped improve notification best practices at scale. In all the above work, locating accurate contact information remains the major challenge regardless of the selected communication channels. On the other hand, finders in [33, 49] collaborated with a medium-sized ISP to explore the effectiveness of email and specialised ISP notification systems against different vulnerabilities, which serves as an example of ISP notification with direct intervention for customers.

Vulnerability disclosure to vendors has received promising best practices as implemented in [3, 42, 44]. Boucher and Anderson [3] have provided detailed guidance on a CVE application, CERT/CC VINCE notification, public press disclosure, and constraints for academic publication. However, despite the stakeholders presenting bug bounty programs and disclosure policies, they encountered slow responses and time constraints in their disclosure operations and in academic publication scheduling. On the other hand, the researchers in the two disclosure operations [42, 44] tackled different software vulnerabilities but encountered the challenge of inviting the same stakeholder to participate in CVD. Lastly, the two operations in [2, 29] cover both vulnerability disclosure and notification. Both revealed the long-lasting challenges of transitioning from disclosure to large-scale notification, as well as the need to explore different communication channels and strategies to mitigate vulnerabilities.

4.2 The Comparison of Different Scenarios

To understand how the finders in the selected work adopted and contributed to best practices, we describe and compare their experiences based on the operation stages introduced in Section 3.1. The comparison follows the order of the stages: Pre-Assessment, Communication Channel, Messaging Infrastructure, Disclosure Policy and Message, and Post-Review.

4.2.1 Pre-Assessment. Once a new vulnerability is found, the finder should assess the potential impact, affected parties, and the associated risk. This will determine the best approach to inform stakeholders. All disclosure operations [2, 3, 29, 44] initiated CVD directly to affected vendors with newly found vulnerabilities and later delegated their disclosure to support organisations. Before their disclosure, most finders could identify several vendors with a bug bounty [3] or tracking program [29], and a short list of direct disclosure contact information [2, 3, 29, 42, 44], which ranged from 5 to 14 affected parties due to limited capacity or testing environments. However, most finders faced challenges such as limited contact information of other stakeholders and the potential development to large-scale disclosure. They then consulted national CSIRTs, CERT/CC (VINCE), or NCSC-NL, except for the researchers in [29]. The impact scale later increased dramatically following cooperation with support organisations, as presented in the Post-Review in Table 1. Despite the lack of contact information and prioritisation in the pre-assessment, the finders could still mitigate the challenges with the support organisations.

The notification operations are vulnerability notifications to the stakeholders, with hosts or domains remaining vulnerable after public disclosure. The finders either leveraged existing datasets [6] or performed network scanning [40] to assess the impact scale of the existing vulnerabilities or other security issues. The scale of the notification operations listed in Table 2, unlike that of the disclosure operations in Table 1, is significantly higher, ranging from 10 to 100,000. In the latest operation in [2], the affected IPv4 hosts have reached a million at an internet-wide scale. On the other hand, the two notification operations in [33, 49] involve fewer affected parties because they were conducted in partnership with a medium-sized ISP and its selected customers. In addition, most notifications included more than one vulnerability or security issue of the vulnerable hosts and domains. Hence, the efforts to estimate the exact number of affected parties increased as well.

As a result, the main challenge of vulnerability notification is how to inform diverse stakeholders and handle multi-party scenarios at scale. Based on the initial estimation of the vulnerability characteristic, impact scale, and affected parties, the next step is to choose the most efficient communication channels and messaging infrastructure to inform stakeholders effectively.

4.2.2 Communication Channel. Selecting a communication channel remains critical [48] to vulnerability disclosure and notification. Depending on the pre-assessment, a finder can select single to multiple communication channels to reach out to stakeholders at scale, as shown in Table 3. With the current guidelines for finders to perform large-scale disclosure and particularly notification, there is still no simple decision-making framework for channel selection that accounts for tradeoffs between contact information retrieval and resource capacity across different scenarios, as discussed in all selected publications.

For vulnerability disclosure, CVD as best practice and support mechanisms, such as bug bounty programs, disclosure programs, and disclosure policies, are provided in guidelines and gradually adopted by vendors [47]. However, improvements to best practices across different scenarios are ongoing. Boucher and Anderson [3] identified their initial list of affected vendors through bug bounty programs, direct disclosure contacts, and outsourced vulnerability platforms. Although the researchers succeeded in most bounty programs, they encountered delayed responses or complications in follow-up conversations with outsourced platforms due to the latter's prioritisation of vulnerability reports [47]. They concluded that direct disclosure contacts or vendor bounty programs might be more effective for extended discussions about their disclosure.

In comparison, Moura and Heidemann [29] experienced late responses and delayed mitigation schedules with a well-known vendor through direct contact and a bug tracking system. On the other hand, all the finders in [2, 3, 42, 44], except for Moura and Heidemann [29], delegated notifications to coordinators such as CERT/CC with VINCE and national CSIRTs. Finally, finders in [29, 42] attended conferences and meetings to reach out to vendors or operators in related communities.

For vulnerability notification, finding an optimised channel to inform stakeholders at scale is more challenging than for vulnerability disclosure. The most significant difference is that the coordinator's delegation to national CSIRTs has been reported as ineffective due to the scale that exceeded the coverage of the support organisations, as reported in [25, 37]. The finders in [36, 37, 50] exhausted multiple communication channels shown in Table 3 and empirically studied the effectiveness of each channel, including remediation rates over time and feedback from affected parties. In the three publications, they all concluded that email remains the prominent channel for large-scale notifications, with the best coverage rate despite obvious pitfalls, such as high bounce rates, spam filters, and low recipient awareness. Similarly, the same issues were reported by finders in [9, 25], who used email as their primary channel.

In the recent five years, finders in [6, 35] adopted best practices derived from the previous work and narrowed their selection to email with a website as support. On the contrary, Maass et al. [26] selected post and email for effectiveness comparison and provided phone support to notified parties. Given the considerable manual effort required to retrieve email and post contact information, they concluded that the post could be an effective but costly channel in their nation. Eventually, the finders in [6, 26, 35] all confirmed that the aforementioned issues of notifying stakeholders via email remain.

Across all selected publications, email is the most widely used communication channel for initial messages and follow-up discussions with vendors and end-users. The main challenges with email remain contact retrieval and low delivery rates. The finders scripted database queries using WHOIS either with a purchased database, as in [36, 37], or with an online query service, as in [6], to retrieve email contact information. Although the finders in [6, 36] mentioned RDAP as a potential alternative, none attempted it. However, Fernandez et al. [12] reported that RDAP has not kept pace with the coverage of the existing WHOIS database, despite the protocol being introduced to improve contact sharing. Moreover, Maass et al. [26] confirmed that manual contact retrieval, at best, does not guarantee accurate contact information either. Aside from the contact retrieval issue, email is constantly challenged by high bounce rate [9, 25], spam filter [6, 36], and low incentives of recipients to read messages [50]. This eventually results in a low delivery rate to the stakeholders.

Nevertheless, a dedicated scalable notification system with a notification partner can be a solution. Finders in [33, 49] teamed up with an ISP and achieved better results with the ISP's walled-garden notifications than email. On the other hand, Beitis and Vanhoef [2] partnered with The Shadowserver Foundation, which actively measures the mitigation progress and notifies its subscribers at scale. Their notification is a new attempt at a communication channel, and the outcome seems promising.

4.2.3 Messaging Infrastructure. A finder will build their messaging infrastructure or delegate it to support organisations based on the selected communication channel. The common practice of outgoing email and website is using a registered individual or disclosure-specific account with the domain name of the finder's organisation to increase the delivery rate and trust of recipients, as implemented in [6, 25, 36, 37]. The email account selection of each selected publication is presented in Table 3. The dedicated account practice is to address the email notification challenge of spam and phishing mitigation. Due to the increasing volume of spam and phishing messages, more and more measures are in place to prevent their delivery. In the early 2000s, it was possible to spin up a mail server and immediately send email notifications to thousands of recipients. These days, there are many different standards associated with sending out emails (SPF, DKIM, DMARC, etc.), as mentioned in [25], and the reputation of the mail server is taken into account before an email is delivered.

Although a finder or its organisation can maintain its messaging infrastructure, it is also reported in [6] that the finders' email infrastructure is partially outsourced to Microsoft Exchange. Their mail server imposes several restrictions, such as limited account control and an email-sending rate limit, which prevent sending large amounts of messages efficiently [28]. This makes it harder for finders to perform large-scale notification operations using email. In addition, this resonates with the trend of vendors outsourcing vulnerability platforms reported in [3], which limits the sender's reachability to the responsible parties, the clients of the outsourced platforms.

There are disclosure and notification operations delegating the messaging to national CSIRTs (such as NCSC-NL), CERT/CC with VINCE [5], and The Shadowserver Foundation. These are listed as Coordinator Delegation and The Shadowserver Foundation in the Table 3. While VINCE serves as a disclosure database with notifications to vendors, as adopted in [3, 44], national CSIRTs have mailing lists or websites to inform vendors and end-users, as noted in [2, 25, 37, 42]. In particular, The Shadowserver Foundation has become a newer notification delegation option with the organisation's own messaging infrastructure to inform the stakeholders, including vendors and end-users, such as ISPs, domain owners, and network operators [34]. Besides, if a finder works with an organisation that has dedicated communication channels, such as an ISP notifying its customers via its internal notification and management tool [33, 49], there can be a more efficient way to trigger action from the recipients. Furthermore, finders in [6, 25, 50] observed that stakeholders, such as domain name owners and cloud providers, used their notification systems to prompt their customers to take action.

4.2.4 Disclosure Policy and Message. The formulation of the disclosure message can determine recipients' attention and whether they will respond and take action [48]. The tradeoffs of length and details of content, such as remediation and security suggestions, are extensively discussed in [6, 25, 35–37, 49, 50], where the authors document not only their disclosure message templates but also the feedback from stakeholders.

Finders in [36, 37, 50] provided brief vulnerability information, affected systems, and disclosure purposes in their initial messages. The messages prompted recipients to either respond to the emails or visit a webpage for more information on vulnerabilities and remediations, using the automatically generated token from each message provided by the finders. This way, the finders could monitor the response rate with a dedicated web backend. Moreover, the finders can reduce the risk of information leaks when recipients are incorrect. Furthermore, a disclosure message could be embedded with HTML content, such as the finder's organisation's logo, to further verify whether a recipient loads the whole message. However, due to spam filters and recent email client loading mechanisms, extensive embedded HTML content is no longer a recommended approach, as confirmed in [36, 50]. From the common experiences of earlier [25, 36, 50] to recent publications [6, 35], the finders confirmed that plain text is the suggested option with less distraction and greater distinction from phishing messages.

In large-scale vulnerability notifications, providing information such as domain names, IP addresses, ports, and issues found in vulnerable systems is recommended to help stakeholders investigate issues promptly, as documented in stakeholders' feedback in [6, 25, 36, 50]. However, receiving parties may have specific mail filters for incoming messages that strictly limit message length and attachment, as reported in [6]. This could hinder the case of cloud providers or domain owners with a large number of vulnerable systems running for their clients. Such stakeholders may decline the messages containing longer vulnerable host lists or any attachments. This eventually makes the composition of a message more challenging for finders [6, 35, 50].

Organisations present support for CVD by publishing a disclosure policy on their websites, using security.txt [18], or providing bounty or tracking programs [3, 29, 43]. Boucher and Anderson [3] observed vendors outsourcing their programs to third-party platforms that reveal their own policies and prioritisation on vulnerability selections. They noted that the policies and preferences of the outsourced platforms may limit the incentives and the delivery of direct messages to the responsible parties. In the vulnerability notification, Chen et al. [6] observed that most stakeholders still do not provide security.txt or equivalent information in their responses, but provide privacy policies that do not necessarily help the disclosure process. The implementation of the disclosure policies among stakeholders still requires attention.

Google Project Zero [19] began with an outgoing vulnerability disclosure policy that outlined the timelines for disclosing vulnerabilities. On the one hand, this has pressured vendors to work on mitigation and publish it within 90 days of the timeline. After several years, the 90-day deadline has become an accepted practice. Academic researchers have also started using outgoing vulnerability disclosure policies [32], which help build trust between stakeholders and coordinate disclosure operations. The outgoing policy has been implemented in [6, 42], where the finders presented their intended procedures on vulnerability handling, notification frequency, public disclosure schedule, and more [31, 39] in their outgoing messages. This gives the receiving party a brief yet informative message for potential procedures and conversation. Other finders in [29, 35, 36, 49, 50] did not document a dedicated disclosure policy, yet did provide equivalent information and contact methods in their messages for coordinating or exempting from the notifications. As for the finders that chose coordinator delegation, ISP intervention, and The Shadowserver Foundation, they also followed and presented the disclosure policy from the support organizations as documented in [3, 29, 33, 36, 37, 42–44, 49]. These delegated parties often use existing relations with stakeholders and customers to create a trustworthy communication channel.

4.2.5 Post-Review. The remediation rate of vulnerable systems is a recurring concern in selected publications, as shown in Table 2. It is worth noting that the remediation rates for each selected publication do not directly reflect the efforts and success of the channel selection and messaging infrastructure provided by the finders. The severity of reported vulnerabilities, the impact scale, the affected system, the implementation of mitigations, and the risk management of receiving parties will all contribute to stakeholders' actions and the progress of mitigation. Although 57% of the vulnerable hosts were patched after the email notification in [9], which is relatively higher than other notification operations with remediation rates from lower than 2.25% to up to around 30% [6, 25, 35–37, 50]. Durumeric et al. [9] initiated the notification of the Heartbleed vulnerability 3 weeks after the notable public disclosure. Besides, they had to drop 56% of the detected vulnerable hosts because responsible administrators were likely to have no access to treat the embedded devices [9]. On the contrary, with a dedicated communication channel and trustworthy notification organisation (ISP Intervention) as in [33, 49], the remediation rate can be significantly higher for more than 75%; nonetheless, such a scenario requires selective recipients or extra capacity and won't necessarily fit other end-user notifications [49]. Moreover, not every disclosure operation can measure the remediation rate due to the nature of the vulnerabilities and affected parties [3, 29, 42, 44]. However, Beitis and Vanhoef [2] include both disclosure and notification operations in a measurable impact scale. They did not provide a remediation rate because their work was still in progress. Their result is worth observing in the near future.

The lack of best practices in large-scale vulnerability disclosure, and particularly in large-scale vulnerability notification, for academic researchers and practitioners, has led to the aforementioned struggles and challenges at each stage, as noted in the title of [29], "Vulnerability Disclosure Considered Stressful". The practice gaps motivated the finders in the selected publications to adopt existing guidelines, reflect on real operations, and contribute to best practices. However, the stress and frustration of the finders deserve attention. Whether it is large-scale disclosure or notification, various vulnerability platforms [3], tracking systems [29], and ticketing systems [6, 25, 50] have increased the workload of finders in delivering messages to the responsible parties in complex multi-party scenarios. The inaccuracy of existing abuse and generic contact information has led to false positives in contact retrieval and to information leaks to unintended recipients, as noted in nearly all selected publications that used email. Besides, despite support from national CSIRTs, finders in [42, 44] still encountered situations in which certain stakeholders did not comply with the CVD as best practice in the first place. The feedback from stakeholders was also not always friendly, either in public or private discussions, as reported in [29, 42, 43]. In some instances, to perform timely disclosure to stakeholders, the finders still need to put in extra effort to contact vendors directly despite having a support organisation as coordinator [42] or disclosure programs provided by stakeholders [3].

Furthermore, reviewing and responding to stakeholders can take time and effort for the finders. Given that the ticketing system is a common practice among stakeholders to receive reports, automatic responses in large-scale notifications can result in a large volume of message content to be examined, as documented in [6, 25, 36, 37, 50]. Even though automatic messages may share patterns to be categorised, the mixture of multiple languages in messages [3, 6, 25, 50], unclear stakeholder disclosure policies [6, 35], and stakeholders' communication systems requiring manual efforts to register and input messages [6, 50] may bring higher than expected workload to finders. This can hinder the effectiveness of large-scale disclosure and notification, and more importantly, the incentives of finders, as discussed in [3, 6, 29, 50].

To sum up, we draw insights from our operations stage model, derived from the selected publications and the experiences shared by support organisations and other communities. This also fulfils our third research question presented in Section 1.1. Furthermore, the insights have highlighted current gaps in best practices for large-scale vulnerability disclosure, particularly for vulnerability notification, in the academic security research community.

5 BEST PRACTICES FOR LARGE-SCALE VULNERABILITY DISCLOSURE AND NOTIFICATION

In this section, we examine the gap in current guidelines and propose new best practices for large-scale vulnerability disclosure and notification based on what we have learned from selected academic publications, communities, and support organisations. The main focus is to help the finders from communities that include academic security researchers, practitioners, and ethical hackers. However, the best practices are not limited to finders but also to stakeholders to improve the receiving strategies for handling vulnerability reports. We follow the same structure as in the previous section to examine limitations and opportunities in the different stages of disclosure and notification operations. At each stage, we discuss common pitfalls to avoid, tradeoffs in method selection, and provide suggestions to the finder and other stakeholders for adopting our best practices.

5.1 Pre-Assessment

Understanding the impact scale, vulnerability characteristics, vulnerability disclosure or notification, and potential stakeholders is essential for a finder to set up a disclosure or notification operation. Whether it's vulnerability disclosure or vulnerability notification to vendors or to end-users will lead to different communication channels, messaging infrastructure, disclosure policies, and message content. From what we observed and as revealed in the reflections of selected publications, due to the lack of large-scale vulnerability disclosure and notification best practices for different communities, a finder may not fully understand efforts and tradeoffs to conduct the disclosure during the pre-assessment stage, which are documented in [3, 29]. We list the points below to highlight the common pitfalls and our suggestions:

- Security researchers, ethical hackers, and practitioners may not be aware of the difference between disclosure to vendors or notification to end-users, prioritisation of the contact list, and the coverage of national CSIRT support. They may then face stress or frustration during the disclosure operations and struggle with unexpected challenges, as noted by Moura and Heidemann [29].
- A finder and a receiving stakeholder may have different risk assessment standards, resulting in different definitions of vulnerability severity from both ends [1, 7]. As revealed in [9, 25], the receiving parties may set a lower priority or lack the capacity to mitigate the issues earlier than the finders expect.
- We recommend timely consultation with CSIRTs or equivalent support organisations. This can help a finder comprehend the potential impact of vulnerability and the scale of the notification to stakeholders, as revealed in the timeline by van Hove et al. [42] and suggested by support organisations [8, 13, 15, 21, 31]. Moreover, this can also help a finder estimate potential stakeholder responses and prepare the message handling in advance.

The following common question is which combination of channel, messaging infrastructure, disclosure policy and interaction can be the most effective in each case. These aspects will be discussed in the following stages.

5.2 Communication Channel

Choosing the most effective channels for disclosure or notification is key to reaching out to stakeholders [48]. We provide recommendations for selecting communication channels for large-scale vulnerability disclosure and notification separately. In the case of vulnerability disclosure to vendors, there are already comprehensive guidelines and support mechanisms in place to help finders inform the receiving stakeholders:

- A finder can seek vulnerability disclosure policies or programs of stakeholders [3, 29], which are possibly indicated by stakeholders providing a security.txt [18].
- A finder can also look for bug bounty programs or vulnerability platforms to issue the vulnerability report. However, stakeholders may also have outsourced such channels to third-party platforms such as HackerOne and Bugcrowd [46], and having a legal agreement in disclosure [27].
- A finder can reach out to national CSIRTs to consult on a possible communication channel [43, 44], if the pre-assessment of potential stakeholders or the impact of a vulnerability is unclear. National CSIRTs may provide notification services, coordination support or platforms such as VINCE by CERT/CC [5] to efficiently identify and inform vendors [15, 31].

In the case of vulnerability notification to end-users, finding an effective and efficient communication channel remains a significant challenge despite the current guidelines:

- A finder can first consider selecting known stakeholders with clear contact information [2, 29], then seek the rest of the stakeholders' contact information.
- A finder should then be mindful of the messaging prioritisation and scheduling if the response time from the initial list of stakeholders takes longer than expected [29].
- A finder should be aware that delegation to national CSIRTs mostly works but may not be effective in every scenario, as the interests, notification coverage, and capacity of different national CSIRTs or equivalent support organisations may vary [9, 36, 50].
- A finder should consider reaching out to key communities or platforms that support remediation tracking [3, 29], depending on the affected parties. Phone calls and posts are more of an option when stakeholders suggest them or the regimes have such practices [26, 36].
- We do not recommend using email to send large-scale vulnerability notifications to end-users, even though it is the most commonly used method. Setting up an email infrastructure for large-scale notifications is difficult, given the spam prevention tools currently in use.
- We recommend the use of The Shadowserver Foundation as a notification channel, as implemented by Beitis and Vanhoef [2]. The organisation can perform active scans and has trusting relationships with key stakeholders that support the notification infrastructure. This has the added advantage that end-users are not overwhelmed by notification campaigns from different finders.

Email, as a communication channel for large-scale notifications, faces the additional challenge of finding contact information for systems on the Internet. The accuracy of the contact information presented on web pages is not guaranteed [26]. Finding contact information for IP addresses is notoriously hard. Although WHOIS and RDAP are methods for retrieving contact information, the results from both are often inaccurate [12]. Moreover, the available contact information is usually intended for abuse notifications, not for vulnerability notifications. While the 'security.txt' standard [18] works for websites, there is no such alternative for IP addresses.

5.3 Messaging Infrastructure

The messaging infrastructure will depend on the selected communication channels. Finders should follow the indicated vendor preference for contact method, usually email or messaging platforms such as bug bounty programs, VINCE, or other vulnerability report platforms. It should be noted that the effectiveness of large-scale email notifications leaves much to be desired, as shown in Table 2. If an email infrastructure is still used:

- We recommend that the sending email address be from a known domain name or with an organisation to increase the delivery rate [6, 36, 41].
- A finder should be aware of the implementation of their mail service; with the trend of outsourcing mail services, there may be rate limits and additional policies to review beforehand [6].

- A finder should be aware that recipients may send automatic responses, divert to ticketing systems, internal communication or management systems, request feedback forms [6, 9, 36, 50], or outsourced platforms with different disclosure policies [3, 42].

The scalability of the messaging infrastructure is the biggest challenge when performing large-scale vulnerability notifications. Although organisations such as CERT/CC with VINCE and national CSIRTs with notification systems can help with large-scale vulnerability disclosure, there is still no optimised infrastructure for large-scale vulnerability notifications in the existing guidelines [24]. Nonetheless, aside from the email infrastructure:

- A finder can provide web pages to describe the intention of the disclosure, vulnerability information, remediation, and disclosure policy. This can help reduce the content in email messages and provide a static source for stakeholders to track issues in the long term [6, 36, 43, 50].
- We recommend checking with a support organisation or partner to determine the available messaging infrastructure. National CSIRTs, PSIRTs, ISPs, and more stakeholders have been improving external or internal notification mechanisms [22, 33, 49].
- We recommend The Shadowserver Foundation, with its established infrastructure and experience in vulnerability notifications, as an effective option for network scanning, identifying vulnerable systems, and notifying affected parties and end-users at scale [34].

In addition, it is worth noting that the adoption of bug bounty programs and vulnerability report platforms as accepted best practices among stakeholders has been growing over the years. However, while well-established stakeholders may have mature report-handling policies and support mechanisms for receiving vulnerability reports at scale [3], other stakeholders may still lack the capacity and experience to establish a clear disclosure policy or program in their public information or ticketing systems [6]. Although some regimes with laws and regulations [8, 11, 31] that mandate stakeholders to take action upon notification have successful cases, such as [26], it is not guaranteed that stakeholders will initiate mitigation, particularly in multi-party notifications at the Internet-wide scale with diverse stakeholders involved, as presented in [6, 41]. Moreover, even vendors and other stakeholders may still struggle to notify their clients and build trust with clients and end-users despite having correct contact information [24]. Eventually, it will be easier for finders to perform vulnerability disclosure and notification at scale as adoption of the practices increases.

5.4 Disclosure Policy and Message

Composing a disclosure message is never trivial. Among stakeholders, a finder and affected parties can have different preferences and procedures for handling the messages [41, 47]:

- In vulnerability disclosure, stakeholders with bounty programs or disclosure policies may provide clear instructions or contact information to initiate the disclosure [46]. In contrast, recipients of the vulnerability notification may not provide enough information on how they will handle a disclosure message.
- A finder should know that stakeholders like network operators may prefer extensive information that includes more details and remediation [31, 41], given that the affected party follows a certain time constraint policy on remediation [6]. Meanwhile, stakeholders such as cloud providers or domain owners may forward the message to their clients with limited communication [6, 25].
- A finder should be aware that not every stakeholder would accept longer messages or attachments regarding the mail server filter, ticketing systems, and spam filter [6, 36, 50].
- We recommend that a finder ensure that the initial message remains brief and does not necessarily reveal every detail of the vulnerability in case of an information leak or legal action [39, 41].

In conversations with stakeholders after initiating vulnerability disclosure or notification:

- A finder should consider that large-scale vulnerability disclosure and notification may get manual or automatic responses from various communication systems in different languages. In most cases, we have seen English used in the responses; other languages are used as support [6, 9, 25].
- A finder should be aware that messages in different languages may increase the processing time and cause confusion, especially if a disclosure policy is presented in a non-native language to a finder or a stakeholder.
- A finder may need to put extra effort into reformulating the messages based on the limitations of a disclosure form or editor in the provided text or system when using stakeholders' communication systems [6].

To inform the intended disclosure or notification, it's important to provide the motivation, the mitigation schedule, and, if possible, the legal terms. As we observed in the selected publications, disclosure policies and equivalent documentation are not widely implemented. This has resulted in both finders and receiving parties having inefficient communication:

- A finder should be aware of the stance of stakeholders during disclosure or notification. Not every stakeholder may want to adopt the best practices for different reasons. This has been reported in two disclosure operations [42, 44] with vendors refusing to participate in the remediation schedule despite having a national CSIRT as coordinator. Still, we encourage stakeholders to adopt best practices and join in disclosure or notification operations.
- We recommend that a finder and support organisations establish an outgoing disclosure policy [31, 32, 39], which includes legal terms, disclosure schedules, message templates, or exemptions, as implemented in [6, 35, 42]. This allows recipients to understand the disclosure procedure from the finder's perspective and protects the finder from unwanted behaviours, such as legal action [27] and public criticism [42].
- We recommend that the stakeholders include a disclosure policy[21, 31, 39] or security.txt[18] in their disclosure programs or responses. This helps a finder initiate disclosure or notification with stakeholders with better preparation.

5.5 Post-Review

With the fast-changing network landscape and the growing number of vulnerabilities, it is important to keep best practices up to date and mitigate security issues in a timely manner. We have seen publications focusing on the exploits, attacks, and network traffic before and after the public disclosure. However, documentation on large-scale vulnerability disclosure and notification is relatively scarce in certain communities, as noted in [6, 29, 36]. Nevertheless, we have seen that finders in [2, 6, 35] could learn from prior best practices and contribute to large-scale notification. Handling disclosure or notification procedures can be a nerve-racking trial, as revealed in nearly all selected publications. It is crucial to have evolving best practices to help a finder prepare for challenges and ease the stress during operations:

- We recommend that security researchers, ethical hackers, practitioners, and other finders document their disclosure or notification operations, including their experiences, considerations, and outcomes.
- We recommend that finders review the impact before and after vulnerability disclosure or notification, which can help finders reflect on their progress and improve best practices.
- A finder can record the challenges and stress encountered during a disclosure or notification operation. The tradeoffs and considerations in different scenarios are also worth recording.
- A finder can provide a timeline as a figure, like in [3, 9, 29, 36, 37, 42] or as text, like in [6, 35], for each disclosure or notification stage, which can help understand the development and its impact over time.
- A finder can track the number of remaining vulnerable systems before and after disclosure or notification, if the vulnerable systems can be traced through network scanning or stakeholder engagement.

- A finder can track the mitigation progress from weeks [9], months [6, 49], or more than a year [35] if the situation permits. The tracking update can be presented as a webpage [4, 44] , experience reports [29], or follow-up academic publications [3, 9, 35, 36, 45] to present disclosure or notification updates.

Last but not least, as we have observed in the academic security research community, academic researchers as finders may show different preferences and limitations compared to bug bounty hunters or vendors as finders. Researchers may need to handle the publication cycle in addition to the disclosure or notification procedure, which adds extra time constraints and stress to their work. Further, we have observed that not every academic researcher has the capacity and prior experience to conduct disclosure or notification operations and finish the documentation within the academic publication period. We believe such a situation can arise in different forms for finders in other communities. With the increasing adoption of support mechanisms such as disclosure and bug bounty programs, the stress of conducting vulnerability notifications at scale can be mitigated, but not significantly reduced, as discussed in Section 4.2.5. This is the main motivation for proposing our best practices to help finders while considering scalability, effective communication, and finder protection.

We have answered our final research question presented in Section 1.1. However, more perspectives on vulnerability disclosure and notification from different finders and stakeholders also require attention. We hope more finders and other stakeholders can benefit from our recommendations and contribute to the security communities by documenting their own experiences and improving best practices.

6 CONCLUSION

The practice of doing academic research on vulnerabilities is growing in popularity. Even though we have best practices for vulnerability disclosure, they require more attention when scaled up. We note a difference in practice between vulnerability disclosure and vulnerability notification, particularly regarding the stakeholders involved. We have analysed trends in academic work and the security community, and proposed new best practices to bridge the gap between existing guidelines and the limitations of operational practice. We believe that our best practices give security researchers, ethical hackers, and practitioners a clear direction to inform stakeholders at scale with less friction. With our suggestions, stakeholders can prepare for the delivery of disclosure or notification messages and for the implementation of mitigation plans. We encourage all stakeholders, including finders, vendors, and end-users, to adopt best practices and document and publish their experiences to improve future vulnerability disclosure and notification operations.

7 ACKNOWLEDGMENTS

This work is supported by the INTERSECT project [23], Grant No. NWA 1160.18.301, funded by the Netherlands Organisation for Scientific Research (NWO). We appreciate the active discussions with the FIRST Vulnerability Coordination Special Interest Group and our colleague, Roland van Rijswijk-Deij, at the University of Twente, which helped us initiate our best practices proposal. We are grateful to the editor and, especially, the reviewers of the ACM journal Digital Threat: Research and Practice, who provided prompt and constructive feedback that helped us refine our content structure and sharpen our research topic. Last but not least, we want to express our appreciation to all the stakeholders who joined us in our earlier notification operation and provided insights that inspired us to compile this best practice.

REFERENCES

[1] Badlock. 2024. https://en.wikipedia.org/w/index.php?title=Badlock&oldid=1206658510

[2] Angelos Beitis and Mathy Vanhoef. 2025. Haunted by Legacy: Discovering and Exploiting Vulnerable Tunnelling Hosts. In *34th USENIX Security Symposium (USENIX Security 25)*. 7135–7152. https://www.usenix.org/conference/usenixsecurity25/presentation/beitis

[3] Nicholas Boucher and Ross Anderson. 2022. Talking Trojan: Analyzing an Industry-Wide Disclosure. In *Proceedings of the 2022 ACM Workshop on Software Supply Chain Offensive Research and Ecosystem Defenses (SCORED'22)*. Association for Computing Machinery, New York, NY, USA, 83–92. doi:10.1145/3560835.3564555

[4] Nicholas Boucher and Ross Anderson. 2023. Trojan Source Attacks. https://trojansource.codes/

[5] CERT/CC. 2023. Vulnerability Information and Coordination Environment (VINCE). https://www.kb.cert.org/vince

[6] Ting-Han Chen, Carlotta Tagliaro, Martina Lindorfer, Kevin Borgolte, and Jeroen Van Der Ham-De Vos. 2024. Are You Sure You Want To Do Coordinated Vulnerability Disclosure?. In *2024 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. 307–314. doi:10.1109/EuroSPW61312.2024.00039

[7] Lucian Constantin. 2016. Hyped-up Microsoft, Samba "Badlock" Flaw Isn't Critical, but Serious Enough. https://www.pcworld.com/article/420541/microsoft-samba-badlock-flaw-not-critical-but-serious-enough.html

[8] Cybersecurity and Infrastructure Security Agency (CISA). 2020. BOD 20-01: Develop and Publish a Vulnerability Disclosure Policy. https://www.cisa.gov/news-events/directives/bod-20-01-develop-and-publish-vulnerability-disclosure-policy

[9] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, and J. Alex Halderman. 2014. The Matter of Heartbleed. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, Vancouver BC Canada, 475–488. doi:10.1145/2663716.2663755

[10] Dutch Institute for Vulnerability Disclosure (DIVD). 2025. https://www.divd.nl/

[11] European Commission. 2022. Cyber Resilience Act (CRA). https://digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act

[12] Simon Fernandez, Olivier Hureau, Andrzej Duda, and Maciej Korczynski. 2024. WHOIS Right? An Analysis of WHOIS and RDAP Consistency. In *Proceedings of the 16th Passive and Active Measurement Conference (PAM)*. Springer. doi:10.1007/978-3-031-56249-5_9

[13] Organisation for Economic Co-Operation and Development (OECD). 2021. Understanding the digital security of products. *OECD Digital Economy Papers* (Feb. 2021). doi:10.1787/abea0b69-en

[14] Forum of Incident Response and Security Teams (FIRST). 2019. CSIRT Services Framework. https://www.first.org/standards/frameworks/csirts/csirt_services_framework

[15] Forum of Incident Response and Security Teams (FIRST). 2020. Guidelines and Practices for Multi-Party Vulnerability Coordination and Disclosure. https://www.first.org/global/sigs/vulnerability-coordination/multiparty/guidelines-v1.1

[16] Forum of Incident Response and Security Teams (FIRST). 2020. PSIRT Services Framework. https://www.first.org/standards/frameworks/psirts/psirt_services_framework

[17] Forum of Incident Response and Security Teams (FIRST). 2020. Vulnerability Coordination SIG. https://www.first.org/global/sigs/vulnerability-coordination/

[18] Edwin Foudil and Yakov Shafranovich. 2022. *A File Format to Aid in Security Vulnerability Disclosure*. Technical Report 9116. RFC Editor. doi:10.17487/RFC9116

[19] Google. 2019. Project Zero: Vulnerability Disclosure FAQ. https://googleprojectzero.blogspot.com/p/vulnerability-disclosure-faq.html

[20] Allen D. Householder and Jonathan Spring. 2022. Are We Skillful or Just Lucky? Interpreting the Possible Histories of Vulnerability Disclosures. *Digital Threats: Research and Practice* 3, 4 (Feb. 2022). doi:10.1145/3477431

[21] Allen D Householder, Garret Wassermann, Art Manion, and Chris King. 2017. The CERT Guide to Coordinated Vulnerability Disclosure. doi:10.1184/R1/12367340.v1

[22] Information-technology Promotion Agency, Japan (IPA). 2024. Information Security Early Warning Partnership | Enhancing Information Security. https://www.ipa.go.jp/en/security/vulnerabilities/partnership.html

[23] INTERSECT. 2020. Towards an Internet of Secure Things. https://intersct.nl/

[24] Rob Knake. 2025. Improving Private Sector Cyber Victim Notification and Support. https://securityandtechnology.org/virtual-library/report/improving-private-sector-cyber-victim-notification-and-support/

[25] Frank Li, Zakir Durumeric, Jakub Czyz, Mohammad Karami, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. 2016. You've Got Vulnerability: Exploring Effective Vulnerability Notifications. In *Proceedings of the 25th USENIX Security Symposium (USENIX Security)*. USENIX Association. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/li

[26] Max Maass, Alina Stöver, Henning Pridöhl, Sebastian Bretthauer, Dominik Herrmann, Matthias Hollick, and Indra Spiecker. [n. d.]. Effective Notification Campaigns on the Web: A Matter of Trust, Framing, and Support. In *30th USENIX Security Symposium (USENIX Security 21)*. 2489–2506. https://www.usenix.org/conference/usenixsecurity21/presentation/maass

[27] Kevin Macnish and Jeroen van der Ham. 2020. Ethics in Cybersecurity Research and Practice. *Technology in Society* 63 (Nov. 2020), 101382. doi:10.1016/j.techsoc.2020.101382

[28] Microsoft. 2023. Exchange Online Limits - Service Descriptions. https://learn.microsoft.com/en-us/office365/servicedescriptions/exchange-online-service-description/exchange-online-limits

[29] Giovane C M Moura and John Heidemann. 2023. Vulnerability Disclosure Considered Stressful. *ACM SIGCOMM Computer Communication Review* 53, 2 (July 2023). doi:10.1145/3610381.3610383

[30] Asuka Nakajima, Takuya Watanabe, Eitaro Shioji, Mitsuaki Akiyama, and Maverick Woo. 2019. A Pilot Study on Consumer IoT Device Vulnerability Disclosure and Patch Release in Japan and the United States. In *Proceedings of the 14th ACM Asia Conference on Computer*

and Communications Security (ASIACCS). Association for Computing Machinery (ACM). doi:10.1145/3321705.3329849

[31] National Cyber Security Centre of The Netherlands (NCSC-NL). 2018. Coordinated Vulnerability Disclosure: the Guideline. https://english.ncsc.nl/publications/publications/2019/juni/01/coordinated-vulnerability-disclosure-the-guideline

[32] Dennis Reidsma, Jeroen van der Ham, and Andrea Continella. 2023. Operationalizing Cybersecurity Research Ethics Review: From Principles and Guidelines to Practice. In *Proceedings of the 2nd Workshop on Ethics in Computer Security (EthiCS)*. Internet Society. doi:10.14722/ethics.2023.237352

[33] Elsa Rodríguez, Susanne Verstegen, Arman Noroozian, Daisuke Inoue, Takahiro Kasama, Michel van Eeten, and Carlos H Gañán. 2021. User Compliance and Remediation Success after IoT Malware Notifications. *Journal of Cybersecurity* 7, 1 (Jan. 2021), tyab015. doi:10.1093/cybsec/tyab015

[34] The Shadowserver Foundation. 2025. https://www.shadowserver.org/

[35] Giada Stivala, Gianluca De Stefano, Andrea Mengascini, Mariano Graziano, and Giancarlo Pellegrino. 2024. Uncovering the Role of Support Infrastructure in Clickbait PDF Campaigns. In *2024 IEEE 9th European Symposium on Security and Privacy (EuroS&P)*. 155–172. doi:10.1109/EuroSP60621.2024.00017

[36] Ben Stock, Giancarlo Pellegrino, Frank Li, Michael Backes, and Christian Rossow. 2018. Didn't You Hear Me? - Towards More Successful Web Vulnerability Notifications. In *Proceedings of the 25th Network and Distributed System Security Symposium (NDSS)*. Internet Society. doi:10.14722/ndss.2018.23171

[37] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. 2016. Hey, You Have a Problem: On the Feasibility of Large-Scale Web Vulnerability Notification. In *Proceedings of the 25th USENIX Security Symposium (USENIX Security)*. USENIX Association. https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/stock

[38] The Shadowserver Foundation. 2025. Open IP-Tunnel Report. https://www.shadowserver.org/what-we-do/network-reporting/open-ip-tunnel-report/

[39] Jeroen van der Ham, Andrea Continella, Petri de Willigen, and Dennis Reidsma. 2023. University of Twente Policy for Coordinated Vulnerability Disclosure in Research. https://www.utwente.nl/en/service-portal/research-support/procedures-facilities/coordinated-vulnerability-disclosure-policy-for-research

[40] Jeroen van der Ham and Roland van Rijswijk-Deij. 2017. Ethics and Internet Measurements. *Journal of Cyber Security and Mobility* 5, 4 (Oct. 2017), 287–308. doi:10.13052/jcsm2245-1439.543

[41] Max van der Horst. 2023. Global Vulnerability Vigilance: Timely Disaster Notification using Internet-Scale Coordinated Vulnerability Disclosure. https://scripties.uba.uva.nl/search?id=record_54279

[42] Koen Van Hove, Jeroen van der Ham-de Vos, and Roland van Rijswijk-Deij. 2023. rpkiller: Threat Analysis of the BGP Resource Public Key Infrastructure. *Digital Threats: Research and Practice* 4, 4 (Dec. 2023), 1–24. doi:10.1145/3617182

[43] Mathy Vanhoef. 2017. KRACK Attacks: Breaking WPA2. https://www.krackattacks.com/

[44] Mathy Vanhoef and Frank Piessens. 2017. Key Reinstallation Attacks: Forcing Nonce Reuse in WPA2. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (Dallas Texas USA). ACM, 1313–1328. doi:10.1145/3133956.3134027

[45] Mathy Vanhoef and Frank Piessens. 2018. Release the Kraken: New KRACKs in the 802.11 Standard. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security* (Toronto Canada). ACM, 299–314. doi:10.1145/3243734.3243807

[46] Thomas Walshe and Andrew Simpson. 2020. An Empirical Study of Bug Bounty Programs. In *2020 IEEE 2nd International Workshop on Intelligent Bug Fixing (IBF)* (London, ON, Canada). IEEE, 35–44. doi:10.1109/IBF50092.2020.9034828

[47] Thomas Walshe and Andrew C. Simpson. 2022. Coordinated Vulnerability Disclosure programme effectiveness: Issues and recommendations. *Computers & Security* 123 (Dec. 2022). doi:10.1016/j.cose.2022.102936

[48] Michael S. Wogalter, Dave M. DeJoy, and Kenneth R. Laughery. 1999. Organizing Theoretical Framework: A Consolidated Communication-Human Information Processing (C-HIP) Model. In *Warnings and Risk Communication*. CRC Press, 29–37. doi:10.1201/9780203983836.ch2

[49] Orçun Çetin, Carlos Gañán, Lisette Altena, Samaneh Tajalizadehkhoob, and Michel van Eeten. 2019. Tell Me You Fixed It: Evaluating Vulnerability Notifications via Quarantine Networks. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. 326–339. doi:10.1109/EuroSP.2019.00032

[50] Orçun Çetin, Carlos Gañán, Maciej Korczyński, and Michel van Eeten. 2017. Make Notifications Great Again: Learning How to Notify in the Age of Large-Scale Vulnerability Scanning. In *Proceedings of the 16th Workshop on the Economics of Information Security (WEIS)*. https://infosecon.net/workshop/downloads/2017/pdf/Make_Notifications_Great_Again:_Learning_How_to_Notify_in_the_Age_of_Large-Scale_Vulnerability_Scanning.pdf