TU Delft

Delft University of Technology

Document Version
Final published version

Citation (APA)
Angelova, Y. Y., Hernandez Ganan, C., Smit, J. A., van Wegberg, R. S., & van Eeten, M. J. G. (2026). Do They Get With the Program? Measuring Mitigation in a Solicited Vulnerability Notification Program. In *35th USENIX Security Symposium 2026* (35 ed.) https://www.usenix.org/conference/usenixsecurity26/presentation/angelova

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.



USENIX

THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

Do They Get With the Program? Measuring Mitigation in a Solicited Vulnerability Notification Program

Yana Angelova, Carlos Gañán, Annebel Smit, Rolf van Wegberg, and Michel van Eeten, *Delft University of Technology*

https://www.usenix.org/conference/usenixsecurity26/presentation/angelova

This paper is included in the Proceedings of the 35th USENIX Security Symposium.

August 12–14, 2026 • Baltimore, MD, USA

ISBN 978-1-939133-58-8

Open access to the Proceedings of the 35th USENIX Security Symposium is sponsored by



جامعة الملك عبدالله للعلوم والتقنية

King Abdullah University of Science and Technology

ARTIFACT EVALUATED
USENIX
AVAILABLE

Do They Get With the Program? Measuring Mitigation in a Solicited Vulnerability Notification Program

Yana Angelova
Delft University of Technology
Y.Y.Angelova@tudelft.nl

Carlos Gañán
Delft University of Technology
C.HernandezGanan@tudelft.nl

Annebel Smit
Delft University of Technology
J.A.Smit-1@tudelft.nl

Rolf van Wegberg
Delft University of Technology
R.S.vanWegberg@tudelft.nl

Michel van Eeten
Delft University of Technology
M.J.G.vanEeten@tudelft.nl

Abstract

Attackers rapidly scan for newly-disclosed vulnerabilities across enterprise networks. In response, defenders have been running similar scans in order to notify the affected enterprise about their exposed attack surface as soon as possible. Such initiatives are becoming more institutionalized, e.g., under the EU's NIS2 directive and CISA's Ransomware Vulnerability Warning Pilot. Research on the effectiveness of vulnerability notifications has found disappointing results, where only a minor fraction of issues were fixed. This was blamed on problems in asset attribution, low trust in the sender, reachability issues, and lacking incentives of the recipient to act. A potential solution to all these problems would be a notification program where companies volunteer to sign up, register their assets, and ensure the right contact details. How much better could such a program perform? We provide the first empirical evaluation of a solicited notification program through a collaboration with a governmental Computer Security Incident Response Team (CSIRT) that offers security notifications to enterprises. We first conduct interviews with nearly half of all participating companies ($n = 21$) to understand why they signed up and how they act on the notifications. Next, we quantitatively study the remediation effectiveness of the program via survival analysis. We find that 27% of the security issues being resolved within one day of notification, 40% within one week, and 49% within one month. Over the entire three years of the program, the remediation rate is 75%. These findings already show higher remediation effectiveness compared to previous unsolicited experiments. This suggests that solicited notification programs can overcome challenges of reachability, trust, and motivation. We reflect on the limitations of these findings and their implications for the future.

1 Introduction

Many organizations operate internet-facing assets with known vulnerabilities and misconfigured services. These issues are leading causes of security breaches and have tangible financial consequences, not merely technical ones. Empirical evidence based on cybersecurity insurance claims shows that misconfigurations and vulnerabilities strongly correlate with the probability of a breach [19, 28, 39].

In response to these issues, public and private organizations have been running internet-wide scans for vulnerabilities and misconfigurations in order to notify the asset owners. Various studies have experimented with such notification mechanisms to measure their effectiveness – e.g., [13, 18, 32, 33]. They found highly variable, and often disappointing, remediation rates, where most of the reported assets remained unmitigated. These studies have identified various explanations for the low remediation rates. First, it is difficult to attribute the asset to the appropriate entity, as many asset owners are not visible in WHOIS or other resources [33]. Then there are major issues with reachability. Most notification mechanisms rely on email, and a substantial fraction of all emails bounce or are otherwise undeliverable [5, 17, 33, 44]. The notification might also not be trusted by the recipient [44]. And, last but not least, many studies question the recipient's incentives to mitigate in response to unsolicited notifications [11, 32, 42].

One solution that would, in principle, address all these factors at the same time are solicited notification programs. These are programs where organizations voluntarily sign up (incentive), register their assets (attribution), and establish a working communication channel (reachability) with the organization they know (trust). Solicited notification programs are often run by government entities. Examples include 'Early Warning' [21] of the National Cyber Security Centre in the UK and the 'Ransomware Vulnerability Warning Pilot' [7] of the Cybersecurity & Infrastructure Security Agency (CISA) in the US. There are also services that are somewhat similar, run by non-profits [29] and commercial providers [2, 30].

How much better do these solicited notification programs perform in terms of mitigation? Although the prior work suggests that these programs *should* achieve better results than the unsolicited notifications studied before, to the best of our knowledge, there has not been an empirical validation of whether this expectation holds up – let alone how much better

the remediation rate would be. In other words, we do not know whether solicited programs should be promoted as the best practice. Given how critical the vulnerability notification mechanism is for protecting public and private organizations, establishing evidence-based guidance is needed.

To address this gap, we partnered with a government Computer Security Incident Response Team (CSIRT) operating a solicited notification program for enterprises in the Netherlands. We posed the following research questions: *(i) What are the considerations and experiences of participants in the solicited notification program?* and *(ii) What is the observed remediation behavior in the solicited notification program?*. To answer these questions, we conducted a mixed-methods study combining interviews with nearly half of all participating organizations ($n = 21$), to get a better understanding of the incentives and experience of participating in the program, with data from the notifications and associated scans ($n = 1{,}766$) to assess whether and how quickly organizations act after receiving a notification. Our study covers three years of historical data (2021-2024) on vulnerabilities, misconfigurations, and compromised systems observed in the assets of the participating organizations. Finally, we examine which factors influence mitigation behavior.

In sum, we make the following contributions: (i) We present the first empirical study evaluating remediation behavior within a solicited notification program. We identify several factors driving organizations to participate in the program, including information gathering, external monitoring, as well as curiosity and a willingness to support governmental initiatives. (ii) We find that 27% of all vulnerable systems are mitigated within one day and 49% of mitigations occurring within a month after the notification. Over the course of the program, 88% is mitigated. These rates are substantially higher than those observed in experiments with unsolicited notifications via email. (iii) We discovered different influencing factors. Critical issues are mitigated faster. Compromised systems are mitigated twice as fast as vulnerabilities and misconfigurations. (iv) We propose practical improvements to enhance the effectiveness and value of solicited notification programs, as reported by the participants themselves.

2 Related Work

A growing body of work examines security notifications, measuring effectiveness and driving factors. We first summarize the state of the art in terms of the observed remediation rates and then focus on the explanations for the limited effect.

2.1 Remediation rates

Prior studies typically focused on one or a few security issues and collected scan data themselves. A classic example is the Heartbleed vulnerability [8]: after its discovery, researchers ran internet-wide scans, sent notifications to affected assets, and tracked remediation for seven weeks. Some work instead notified about compromised systems. We distinguish between vulnerabilities/misconfigurations and compromises when surveying remediation rates.

Vulnerabilities & Misconfigurations. Studies on vulnerability notifications report mixed results, further complicated by differing observation windows. We first summarize findings as reported, then identify remediation rates at more standard intervals (one day, one week, two weeks, one month) where possible.

The 2014 Heartbleed experiment by Durumeric et al. [8] used email notifications and observed close to 50% mitigation in 14 days for a high-profile vulnerability with an immediate compromise threat. Several later studies also used email. Li et al. [13] reported only 18% remediation over 14 days for a mix of issues, including exposed industrial control systems and misconfigured IPv6. Cetin et al. [5] likewise found 18 remediation over 19 days. Stock et al. [32, 33] extended the observation window to about one month and found mitigation rates around 25%. Utz et al. [37] notified about various privacy and vulnerability issues, tracking effects for two months and observing rates of roughly 10–15%. Longer-term studies over five to six months by Lone et al. [16], Nosyk et al. [24], and Cetin et al. [43] show similarly modest effects: Lone et al. found that email notifications about source address validation had no significant impact; Nosyk et al. reported around 25% remediation; Cetin et al. observed about 30% improvement relative to a control group.

Several studies have explored channels beyond email—instant messaging, browser alerts, phone calls, and postal mail. Zhang et al. [42] used instant messaging in an eight-day campaign targeting vulnerable campus networks and observed remediation rates around 40%. Maass et al. [17] achieved 49% remediation in one month using postal mail. Maass et al. [18] combined email and letters over roughly two months and reported 57% remediation, while Zeng et al. [41] used browser alerts over a similar period and observed 34%. Sasaki et al. [27] studied phone calls with administrators of vulnerable systems over approximately six months and found that successfully reaching the responsible person led to a remediation rate of 58%. The high rates in Maass et al. [18] also reflect the use of warnings—arguably threats—of potential monetary fines for privacy misconfigurations.

Overall, these experiments span a wide range of mitigation rates. Channels beyond email can substantially improve remediation but are difficult to scale: postal mail is expensive; discovering suitable instant messaging or social media contacts is labor-intensive and often impossible; and legal threats are neither feasible nor desirable in many contexts.

To better compare all these previous studies, we want to put them on an equal footing and extract from those papers the remediation rates at four different time points: 1, 7, 14, and 30 days post-notification. For this purpose, we used either the numbers explicitly discussed in the paper, or we inferred the

number from the survival plots and cumulative distribution function graphs. The one study excluded from this plot is that of Cetin et al. [43] due to the lack of information on the remediation speed. The results are plotted in Figure 1. Note that the lines connecting the dots are there for legibility. They are not the curves from the original studies.

Due to space limitations, we cannot repeat the figure later in the paper. So we include here already the findings from our study. Since the program that we evaluate sent mostly notifications on vulnerabilities and misconfigurations, we focus the comparison on those findings specifically. As can be seen in Figure 1, the program we studied has mitigation well above the comparable email-based experiments. The mitigation rates that are closer to it are all studies using different mechanisms, like postal mail and instant messaging, as well as deploying legal threats. The only exception is the original Heatbleed study (Durumeric 2014). This was a highly-publicized and critical vulnerability that faced large-scale exploitation. The program we evaluated sent out notifications on more mundane issues, yet still achieved similar remediation rates.

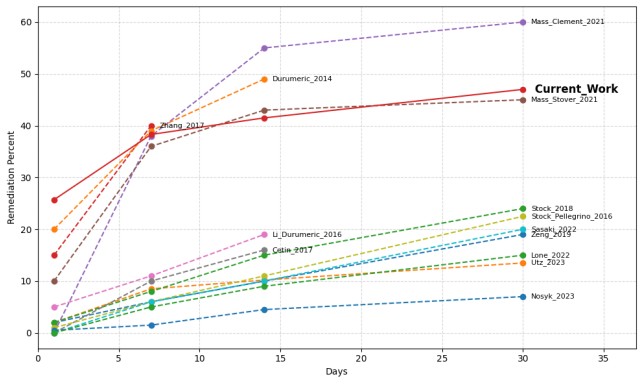


Figure 1: **Remediation rates from prior work standardized at intervals of 1, 7, 14, and 30 days. Lines connecting those dots are not the curves from the original plots.**

Compromised. Some experiments focused on compromised and abused systems. One of the first ever studies in the field of security notifications, by Moore and Vasek [38], showed that 32% of the affected systems were cleaned up within a day. More and Vasek's rates improve significantly once the monitoring period increases to 16 days, and they observe that by then 62% of the systems would be mitigated. This rate is comparable to the one presented by Cetin et al. [44] who also observed between 50 and 60% remediation after 16 days. Finally, Li et al. [14] also reported remediation at 60% but on a much larger time frame of 11 months. In short: compromise issues prompt more and faster remediation compared to vulnerabilities and misconfigurations.

2.2 Influencing factors and challenges

Different factors have been reported to influence the success of security notifications. One main difference in the mitigation rates we have described above concerns the type of security issue: compromised systems are prioritized over vulnerable or misconfigured systems. This pattern suggests that the perceived severity, urgency, and relevance of the issue strongly influence recipients' motivation to act. Compromise-related notices likely convey a more tangible sense of risk.

A recurring concern for low remediation rates is the problem of reachability. Works by Stock et al. [33], Li et al. [13], and Zeng et al. [41] all highlight that the impact of notifications is highly dependent on reaching the recipient. Factors such as outdated, missing, or hard-to-find contact information, and bounced emails hamper notification efforts. Some studies [5, 34, 44] report email bounce rates of up to 50%. This issue was further highlighted by Soussi et al. [31], who concluded that directly emailing the end-owner of affected asset is likely effective, but it is often not scalable due to the complexities involved in identifying those end-owners, especially after the General Data Protection Regulation (GDPR) came into effect in the European Union. [26] showed that missing, outdated, or improperly configured security.txt files, mechanisms intended to facilitate security communication, can impair future vulnerability disclosure. Such observations illustrate the critical need for improved attribution and contact discovery mechanisms.

Beyond reachability, trust in the sender is another key factor affecting notification effectiveness. Hennig et al. [11] found that recipients often distrust unsolicited security alerts, suspecting phishing or scam attempts, which in turn lowers the likelihood of timely remediation. While Cetin et al. [44] found no significant correlation between sender reputation and response rates, other studies, such as those by Li and Durumeric [13], suggest that notifications from trusted or well-known organizations tend to achieve higher response rates than those sent by independent researchers or intermediaries. Several studies have examined the impact of message framing and recipient characteristics on the effectiveness of security notifications. Works by Mass et al. and Utz et al. [18, 37] suggest that threat- or legal-risk framings present stronger responses than privacy-oriented messages. Recipient characteristics also matter: Stover et al. [35] identified personas shaped by technical expertise and organizational capacity, suggesting the need for tailored notification strategies.

A final factor is the effectiveness of the communication channels. Some studies suggest that postal mail and instant messaging sometimes outperform email, while social media and web forms generally underperform [17, 32, 42]. Email is still the default in practice. Any alternative, like postal mail or instant messaging, runs into the problem of scaling, rendering them impractical for large-scale real-world applications.

Academic studies of notification mechanisms have, so far,

focused on ad-hoc, short-lived efforts that sent unsolicited notifications, where alerts are sent without the recipient's prior request or expectation. Such efforts lack follow-up, consistency in communication, and long-term engagement with the affected parties. They stand in contrast to solicited notification programs. Our study focuses on the latter and aims to compare its remediation success to the rates reported for unsolicited, email-based notification efforts.

3 Data & Methodology

This work is built on three main datasets: (*i*) interviews with participating organizations in the solicited security notifications program, to identify their expectations and experiences with the notifications; (*ii*) ticketing data that contains all the notifications sent to the participating organizations; and (*iii*) daily vulnerability scan and abuse data that the notifications are based on, which allows us to track whether an issue is remediated. The latter two sources allow us to measure whether and how fast remediation takes place after the notification.

The national CSIRT for the Netherlands announced the solicited notification program in 2021. (Technically, at that time, the organization initiating the program was a separate government entity. It was later merged with the national CSIRT.) It invited organizations to express interest in joining. From this larger group, it selected 57 organizations ranging from small to large across nine different sectors. They entered the program in September 2021. A table with the demographics of all program participants is available in Appendix E. One of the criteria for participation was that those organizations should not already receive security notifications from other entities providing such data, as recognized by the government (e.g., from other programs at the national CSIRT itself, from sector CSIRTs, or other local information-sharing programs).

The selected organizations registered all their internet-facing assets with the CSIRT and provided a dedicated email address for receiving security notifications. The CSIRT receives daily vulnerability scans and abuse feeds from various sources, mostly other CSIRTs and security non-profits. Each day, they then compare the received data against the registered assets of the participating organizations. If there is a match, it sends a notification email to the affected organizations. If the issue is not resolved and the same asset continues to appear in subsequent scan data, the CSIRT sends a reminder notification after 28 days. It repeats this action every 28 days until the issue is resolved.

3.1 Interviews with program participants

To better understand the considerations associated with participating in the solicited notification program, we set up an interview study with the registered organizations. We use the qualitative data to provide first-hand insights from the participants themselves and contextualize the findings from the analysis of the notifications.

Respondent selection. We contacted all 50 organizations participating in the program as of October 2024 (seven had dropped out earlier), using the email address they registered for notifications. The recruitment email – sent from the author's government-issued CSIRT address – explained the study's purpose and stated we would follow up by phone, including the specific number we would use. Organizations could opt out of the interviews by simply not responding or by replying to the email.

In the following weeks, we called each organization twice and sent one reminder email. Of the 50, twelve never responded, seven declined, and six initially agreed but later disengaged. We scheduled interviews with 25 organizations; four later canceled and did not respond to rescheduling attempts. Ultimately, we conducted 19 interviews, including one joint session with three organizations acting as private CSIRTs for their sector, resulting in 21 respondents—nearly half the participant pool. All interviews were conducted online between April and June 2025. Respondents were directly involved in handling the DTC's notifications.

Interview protocol. The interview protocol was developed through several iterations. We tested it in two pilot interviews. One interview was with a security administrator processing notifications at the authors' institution, and the other with the security officer of a non-profit organization. After the two interviews, minor changes were made to the protocol, resulting in the final version presented in Appendix A. The same protocol was used in all interviews, with the only change being the language in one of the calls, where the participants requested the use of English rather than Dutch.

Interview coding. We recorded all interviews under informed consent. Each recording was fed into a local instance of the AI transcription model Whisper [40] running on a secure server. The transcripts were later checked by one of the authors against the recording and corrected where necessary.

We used ATLAS.ti [3] for coding the transcripts. The lead author selected a random sample of 5 transcripts and used inductive coding to create the initial codebook. Most interview questions were coded by applying an inductively developed set of categories that were mutually exclusive and exhaustive, while a subset of more open-ended responses was analyzed thematically. They then provided the codebook and one of the transcripts to the second coder for independent coding. The two coders discussed the coded document and refined the codebook to remove ambiguity. The lead author then re-coded the five transcripts with the refined codebook again.

The five transcripts were then given to the second coder for independent coding using the refined codebook. Any disagreements were further discussed between the two coders, and the code definitions were refined. For the questions coded using thematic analysis, the two coders independently coded the five transcripts using the refined codebook and then discussed

their findings together to refine the codes and definitions further. We did not calculate Inter-Rater Reliability (IRR), but did track and discuss disagreements as stated above; this is in line with Braun & Clarke, and others, who note that IRR is not usually a measure of quality for codebook-based TA [4, 20]. The lead author then proceeded to code the remaining transcripts with the final codebook, which is presented in Appendix B.

3.2 Security notification data

One of the authors was set up at the CSIRT as an intern. They were provided with access to the ticketing system used for security notifications. It is essential to clarify that we did not perform any notification experiments ourselves, but instead used historical data from the CSIRT's solicited notification program. The system contains all security notifications sent to all participating organizations from October 2021 to October 2024. The processing of this data can be split into two main phases, as illustrated in Figure 2. In each phase, we transform and enrich the notification data, enabling us later to evaluate the mitigation actions in response to CSIRT notifications.

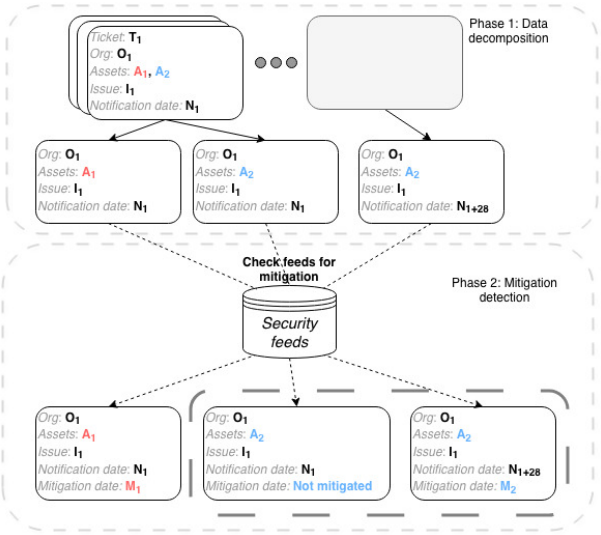


Figure 2: **Representation of the different steps in the notification data processing**

Phase 1: Ticketing data decomposition In the system, each email notification is recorded as a separate ticket. The saved data consists, among other things, of the ticket creation date, the email address of the recipient, associated vulnerable, misconfigured, or compromised assets, the specific type of vulnerability, misconfiguration, or compromise, and all correspondence between the affected organization and the CSIRT regarding the notification, including possible attachments with more information about the issue at hand. We have enriched the notification data by incorporating the severity label as reported in the underlying security feed. Note that the CSIRT does not label the severity in their notification, except for 'critical' issues.

Because the ticketing workflow at the CSIRT groups together affected assets from the same company, security issue, and date into a single ticket, we first had to split these into separate data points. This is necessary since not all affected hosts may exhibit the same mitigation behavior. The resulting dataset after this step consists of data points corresponding to a unique combination of affected asset (A_i), security issue (I_i), and notification date (N_i).

Phase 2: Mitigation date detection To measure the effects of the security notifications, we need to know if the affected assets are mitigated or not. For this purpose, we used the daily security feeds of Shadowserver [29]. Shadowserver provides daily network scan reports for over 100 different types of security issues. It has been the primary source used by the government's CSIRT in its notification process. We focus our analysis fully on notifications that were based on the Shadowserver data, since we have the full daily scan reports for this source, but not for other sources that the CSIRT used. To detect the date of mitigation, we check when, after a notification, an asset is no longer visible in the Shadowserver feeds.

There are limitations to this measurement of mitigation. One common issue concerns asset churn, e.g., changes in ownership. For our study, this has a limited effect since we observe enterprise asset space, which is more stable, and the organizations have self-registered their IPs and can update them whenever changes do occur. Another limitation concerns the assumption that the absence of an asset in the feed means the issue is mitigated. To try to work around this issue, we also track whether it later re-emerges for the same asset, which would trigger a re-notification from the CSIRT. In these cases, we assume the initial disappearance was temporary and should not be counted as mitigation. That being said, not every remediation is observable from these scan results. Affected assets might be remediated by other security controls, such as firewall rules that block exploitation attempts. This means that the vulnerability still exists on the asset, but that the risk is mitigated via other means that are invisible to the scans.

All in all, we have created a final set of 1,766 notification cases spanning three years (2021-2024). Each cases consists of a unique combination of Internet-facing IP address, type of security flaw, and notification dates. The complete list of variables in the dataset is presented in Appendix C.

4 Limitations

Our study has several limitations. First, we evaluated a single national, solicited notification initiative. It represents the type of institutionalized mechanism envisioned under NIS2 and similar policies. These are still emerging, so the value of our analysis lies not in generalizing to all settings, but in establishing an empirical estimate of remediation performance. Vulnerability disclosure and notification programs typically

follow internationally shared guidelines and practices (e.g., FIRST CSIRT Services Framework [10]), resulting in substantial similarity across countries that adopt them, particularly in Western industrialized contexts. That said, future studies are needed to establish a broader geographical base of evidence for effectiveness.

Second, our study would ideally have compared the solicited program to a control group. We study a real-world governmental service, ethical constraints and program terms-of-use prohibit withholding notifications from affected organizations. This means a control group was infeasible. This limits causal inference, but it does not exclude measurement: we benchmark the observed remediation rates against those reported in all prior work. We acknowledge, however, that this comparison is impacted by selection bias. Participants in solicited programs are likely more motivated and more mature than the broader population targeted by non-solicited programs. This means we cannot separate the effect of the program itself from the effect of the differences in recipient populations. Despite these caveats, benchmarking against prior work remains valuable as it situates our findings within the broader literature and provides a point of reference in the absence of a feasible control condition. A causal evaluation with an explicit control group remains an important direction for future research, though such designs run into ethical issues that might not be surmountable.

Third, we observed selection bias in the interview recruitment. Interview participants exhibit faster remediation than non-participants. This self-selection limits the generalizability of qualitative insights to less proactive organizations. However, our quantitative results cover the full dataset of notified organizations, ensuring that the measured remediation rates are not affected by interview participation. There is a second source of selection bias, namely more security-conscious organizations are more likely to sign up for such a program. So our findings are likely an upper-boundary estimate for remediation when generalizing the program to a broader sample of the population.

Fourth, there are limitations in our remediation detection. We infer it from disappearance in Shadowserver feeds, the same data the CSIRT uses operationally. This approach misses certain aspects, as discussed above. We would argue these limitations bias estimates downward, meaning our reported 88% remediation rate is a conservative lower bound. The qualitative findings corroborate this interpretation, with participants consistently reporting action on all issues.

Taken together, these limitations reflect the realities of studying operational CSIRT programs. Rather than calling into question the findings, they clarify the boundaries within which our results should be interpreted and highlight promising directions for future evaluations.

5 Participant perceptions of the program

The aim of the interviews was to gain a deeper understanding of the considerations of organizations for signing up, their experiences, notification processing, and the future improvements they envision for the program.

The 21 organizations we interviewed cover all of the sectors involved in the program except Transport & Logistics. A table with the complete demographic information of our interview respondents can be found in Table 1. In terms of organizational size, the sample comprises 14 large organizations, four medium-sized entities, and three small ones. For categorizing organization size, we use the standard European categories [9] of small ($\leq$49 FTE), medium (50-249 FTE), and large (250+ FTE). We asked the organizations about the size of their IT staff. Only four reported having IT departments comprising more than 50% of their total FTEs. At the other extreme, one organization operated with a single ICT staff member. The remainder had IT departments comprising less than 20% of their total workforce. All interviewees had been directly involved with the program's security notifications and held technical roles, with five in lead management positions. Among the respondents were four CISOs and four security officers, three security analysts, three system administrators, one security engineer, and one system engineer. All but four reported having more than five years of experience in their current positions.

5.1 Incentives to participate

Understanding why these organizations signed up helps to understand their behavior, as well as inform strategies for recruiting future participants, since the program aims to expand in the near future.

We first attempted to determine how the program was introduced in the context of the organization. Was participating a bottom up idea, initiated by security personnel to improve security? Or was it initiated by higher management, for example, as a way to show best practices and compliance? From our interviewees, we found that nearly all respondents (20 out of 21) initiate the sign-up process themselves, so mostly from the bottom up. Four of our interviewees were CISOs. While not truly operational security personnel, they still sit in the middle between the operational personnel and top leadership. In only one case, we saw that the initiative came from higher management: *P11:"[the program] was pointed out by our previous CEO"*.

Next, we aimed to determine the considerations for joining the program. Respondents could mention multiple reasons. Three main themes emerged as reasons. The most often mentioned reason (14 of 21 participants) was receiving **additional security information** from the governmental CSIRT. They valued the potential for new sources of information that the government might have. More generally, they mentioned

Table 1: **The demographics of interview participants**

ID	Position[a]	Experience[b]	Org. size	ICT %[c]
P1	SYE	20 years	Small	50+%
P2	SO	3 years	Large	≤20%
P3	CISO	16 years	Small	50+%
P4	Other	15 years	Small	50+%
P5	SEE	4 years	Large	≤20%
P6	Other	4 years	Medium	≤20%
P7	CISO	6 years	Large	≤20%
P8	CISO	17 years	Large	≤20%
P9	CISO	11 years	Large	≤20%
P10	SYA	5 years	Medium	≤20%
P11	SYA	20 years	Medium	None
P12	Other	6 years	Large	50+%
P13	SYA	19 years	Large	≤20%
P14	SO	7 years	Large	≤20%
P15	SO	26 years	Medium	≤20%
P16	Other	1 year	Medium	≤20%
P17	SO	13 years	Large	≤20%
P18	Other	26 years	Medium	≤20%
P19*	SEA	7 years	Large	≤20%
P20*	SEA	10 years	Large	≤20%
P21*	SEA	6 years	Large	≤20%

[a] Interviewee Position – SYE: System Engineer, SO: Security Officer, SEE: Security Engineer, SYA: System Administrator, SEA: Security Analyst
[b] Reported years of experience on the current position within the org.
[c] Reported proportion of the ICT department size of the organization
* Participants interviewed together

that every piece of extra information is welcome for improving their cybersecurity posture. Most of the organizations reported that you *"can never have too much security information" (P15)* and *"all information sources have added value" (P8)*. Some, like P4 mentioned that they *"... always keep on improving their security [with such information]"*. Some organizations expected the information to be about general and new security trends that the CSIRT observes, and they admitted to being pleasantly surprised that they received tailored security notifications relevant to their own organizations.

Another recurring reason for joining, mentioned by 8 of our interviewees, was the **added value of external monitoring**. Organizations found it helpful to have *"an extra pair of eyes when it comes to security" (P2)* that could help them *"monitor [their attack surface]" (P7)* or use it as a *"supplement ... to learn from" (P5)*. Some see the value in the awareness the program brings, specifically the fact that *"when you are attacked" (P2, P14)* or *"appear in a vulnerability feed" (P13)*, the CSIRT will notify you, as an additional instrument for signaling when problems occur.

Among the **other reasons**, we also see some participants state that they were incentivized because it was a governmental initiative that sounded interesting and free. As one participant put it: *"And since it was free, which is like the ideal price, right?" (P19)*. This suggests that some organizations signed up out of *"curiosity" (P1)* and to *"support such governmental initiatives" (P12)*, without having explicit expectations on how the program would benefit their own organizations directly. Interestingly, four organizations told us that for them these security notifications are a form of *"extra layer of security" (P18)* that, if all else fails, is still there as a form of backup. This 'extra layer' is more abstract than getting operationally relevant security information. They believe that participating in the program is valuable, but it does not serve an operational role in their security practices; rather, it serves as a safeguard to ensure that what they are already doing is correct.

We observed that interviewees with less work experience were more likely to identify the operational benefits of additional information and external monitoring, while the "other reasons" were more often mentioned by experienced respondents. This suggests that technical details are more useful for less experienced security personnel, while mature employees usually value such initiatives for non-technical reasons.

Remarkably, there is no clear connection between the reasons and the organization's size or the size of its ICT department, which could suggest that the considerations are more driven by specific professionals, rather than a function of organizational properties. Another interesting detail is that organizations with in-house vulnerability scanning solutions do not mention the value of external scanning as a reason for joining the program.

5.2 Participant experiences

How did participant expectations evolve during the program? Almost all interviewees were very positive regarding their experiences, mentioning that their initial expectations were met. One respondent, who had experienced asset registration issues, emphasized not ending on a negative note to avoid giving the wrong impression: *P15: "And yes, that may sound a bit negative, but I am happy with how the CSIRT has worked, you know. So I don't want to end on a negative note."*.

Most respondents (14 of 21) told us that their organizations had received very few notifications in the past three years. Only two interviewees mentioned having uncertainty due to the lack of communication from the CSIRT in the beginning: *"So you're left wondering, what is the CSIRT actually doing?" (P4)*. This could be improved by more regular communication, even in the absence of specific notifications to the organization. Nevertheless, most interviewees do not perceive this as a negative feature of the program. On the contrary, six interviewees actually mentioned that they take that as a good sign regarding their security posture. As P10 put it: *"What's also nice is that I don't get many messages about [security issues]. So that means we're doing a good job"*.

A surprising finding was that no one mentioned receiving false positive notifications. There have been cases where the

notifications were not directly relevant to the recipient. For example, the asset no longer belonged to the organization, and they had forgotten to update their registration with the CSIRT. Or the asset was managed by a client of the organization, and they had no oversight of what the client was running. In other cases, the vulnerability that they were notified about had already been mitigated, but this was not visible to outside scanning, meaning that the assets were still shown as vulnerable. For all of these cases, participants in the program would like to have an option of a feedback loop, where they can report alternative mitigation or indicate that the information was received but falls outside of their responsibility.

5.3 Notification processing & prioritization

Participants described a broadly similar workflow for processing notifications: verifying reproducibility, assessing remediation options, and mitigating the issue. As one summarized, *"Well, we open the drawer, get a pitchfork and a torch, and then start" (P21)*. Interviewees emphasized that for smaller organizations with limited security staff, identifying and validating issues can be time-consuming without clear explanations and guidance in the notification. By contrast, experienced practitioners in larger organizations reported that a concise notice often suffices, though even there, additional detail may be needed when escalating issues to management or affected clients. All participants indicated that they address all issues they are notified about, except when these fall outside the organization's purview, such as assets managed by clients or third parties that are no longer owned by the recipient organization but have not yet been de-registered with the CSIRT.

In our interviews, we included an interactive component, where the participants were presented with five sample notification emails and asked to assess them. The goal was to directly observe recipients' decision-making processes on different issues. We could then later compare those reported behaviors with the observed remediation behavior in the scan data (Section 6). The sample included one notice about a compromised system, two about vulnerabilities, including a vulnerable Microsoft Exchange Server, a well-known target for adversaries [6], and two about misconfigurations. All sample notifications are in Appendix D.

Of our 21 respondents, 19 were willing to prioritize the notifications. The remaining two said they would check the CVSS score or whether the issue was prioritized by the national CSIRT. We asked the 19 interviewees to think out loud while processing the notifications and to rank them as if all were received simultaneously. Given that such volumes are atypical, one interviewee joked that if this would happen, *"I will directly go on vacation" (P13)*, underscoring how extraordinary such a scenario would be.

The issues that our respondents prioritized as requiring immediate attention are most often those concerning an already compromised systems or a well-known vulnerability that is actively exploited. From the 19 respondents who shared their prioritization with us, 11 put the compromise with the highest priority, while five gave urgency to the vulnerability concerning an Exchange Server. Another prioritization strategy mentioned is based on the severity assessment of external entities, such as the National CSIRT (NCSC) [22] or the National Vulnerability Database (NVD)'s Common Vulnerability Scoring System (CVSS) score [25]. Although, in essence, this strategy also relies on the type of security issue and its perceived severity, the main difference with the rest is that it removes assessment and classification by the recipient and instead relies on the judgment of external entities. As P11 put it: *"I will check the NCSC RSS feed to see if there's anything there. Which one is the most important... the most dangerous."* Another respondent noted: *"Of course, if issues with a severity rating of medium or higher are identified, we will assess whether they can be addressed as quickly as possible." (P4)*. These responses suggest that some organizations, often smaller ones, benefit from clearer prioritization guidance, a function not currently provided by the CSIRT.

5.4 Future improvements to the program

We ended with a question about program improvements our interviewees would like to see. Here, three main themes emerged most often. By far the most requested improvement is the development of a **dashboard or portal**, where organizations can see historic notifications, manage their assets, indicate preferences for what notifications to receive, and report mitigation actions that are otherwise invisible to external scans. They mentioned that the current mechanism for most of these actions is via email, which is very tedious and often discourages them from contacting the CSIRT. As one interviewee mentioned: *"That's a very good addition, because I've always found it burdensome to get something new added. So, a separate portal to view notifications. I think that's a good one."(P5)* This also explains why 12 participants reported not having updated their assets since the initial sign-up, even though eight of them reported that there have been changes to their underlying asset list. Keeping asset registrations up to date is not only valuable for the participants but also essential for the notification program as a whole, as it is a prerequisite for the proper functioning of the notifications.

Another theme we identified concerns the communication around the program. Seven interviewees stated that they would like to see **improved communication** from the CSIRT. Several noted that the current notifications and security updates are often written for audiences with a strong technical background. Some participants said they sometimes need to share these notifications with clients who are less technically capable. One participant also speculated how this would work when the program expands and includes less mature organizations: *"If you need to explain this to the bakery around the corner here, do they know what DNS is and a TXT record? So*

they probably need to translate it to their IT supplier, or their nephew who's doing their computer stuff." (P20). Others mentioned being left in doubt about the program's ongoing efforts due to the lack of notifications they received over a prolonged period. As P4 mentioned: *"So you are left wondering, what is the CSIRT actually doing?"*. Overall, participants expressed a desire for more consistent communication that is tailored to the varying levels of technical expertise within the program.

A final recurring theme in our interviews was the desire for **stronger community engagement**. Apart from the email notifications, participants in the solicited notification program have had access to a forum facilitated by the CSIRT. There, they can exchange ideas, news, and information directly with one another. They have also been invited to dedicated "CSIRT Days" once a year, where they can meet the team behind the program, get more information, and provide feedback regarding the future of the notification program. Seven participants reported that they already make use of the Forum and CSIRT Days, and emphasized the added value of directly connecting with peers to exchange ideas and information. At the same time, five interviewees expressed a wish for even more community-building and support for information sharing, noting that the CSIRT is uniquely positioned to facilitate this, given its cross-sector vantage point. Several larger organizations also suggested that a form of "reverse" information flow could be valuable. Participants 19, 20, and 21 shared past examples where they had informed the CSIRT of other vulnerable organizations so that the CSIRT could notify them directly. In this way, program participants could help strengthen the CSIRT's efforts by sharing insight that might otherwise remain unknown to the government.

6 Measuring the impact of notifications

In this section, we present the results of our quantitative analysis of security notifications. Our goal is to measure the speed of remediation and potential influencing factors.

6.1 Distribution of notifications

The notifications span three years, from October 2021 to October 2024. We have identified 1,766 data points, i.e., notification cases. To provide context for the measurement of remediation, we first explore who received the notifications.

From Table 2, we see that our notification data primarily concerns small and large organizations, with medium-sized ones accounting for only 6% of notified affected assets. Most notifications (54%) go to small organizations. This is disproportionate compared to the distribution of the 50 participants in the program, where just 16% are small organizations (19% medium, 65% large). Given that smaller organizations also have fewer assets on average, this suggests that they have a higher frequency of security issues compared to larger ones.

Table 2: **Summary of notification cases** ($N = 1{,}766$)**.**

	#	%
Org. size		
Small / Medium / Large	960 / 108 / 698	54 / 6 / 40
Issue type		
Misconf. / Vuln. / Comp.	1,013 / 685 / 68	57 / 39 / 4
Severity		
Low / Med. / High / Crit.	89 / 811 / 508 / 358	5 / 46 / 29 / 20
Re-notified		
Yes / No	778 / 988	44 / 56
Mitigated		
Yes / Unknown	1,330 / 436	75 / 25
Interviewed		
Yes / No	210 / 1,556	12 / 88

The organizations self-reported which sector they belong to. Most of the notifications were sent to organizations in the ICT sector (68%). The remainder is distributed across many other sectors, the full distribution can be seen in Appendix F. This distribution is somewhat surprising, given that the ICT sector is not the sector with the most internet-facing assets registered with the CSIRT. This discrepancy could be attributed to the fact that some of the organizations in the ICT sector act as intermediaries for third parties, expanding and diversifying the surface of responsible entities and systems.

In terms of the severity of security issues, as classified by ShadowServer, only a few notifications were about 'low' severity issues. This reflects the policy of the governmental CSIRT to focus on higher-severity issues. They want to avoid overwhelming organizations with notifications that might drown out more critical issues. Most of the notifications from the CSIRT concern issues with misconfigured or vulnerable systems, while only 4% are about compromised assets. In the cases where remediation of the security flaw does not happen, the CSIRT sends reminder notifications. In our data, approximately 44% of all notification cases have one or more reminder emails, indicating that they were not resolved within 28 days, as far as the CSIRT could tell. With the number of reminder emails sent ranging from 1 to 38, one might wonder to what extent such re-notification actually affects mitigating the underlying issue.

We will revisit these properties of the notification cases in the next section to see if they correlate with remediation.

6.2 Remediation rate & influencing factors

We now turn to our research question: *What is the observed remediation behavior in the solicited notification program?* Over the roughly three-year program, we observed remediation in 75% of all cases. Because this raw rate depends on program duration, we apply survival analysis with right-censoring, the standard approach for such processes. Figure 3 shows the Kaplan–Meier survival curve. Survival analysis treats every notified asset as a case that needs to be reme-

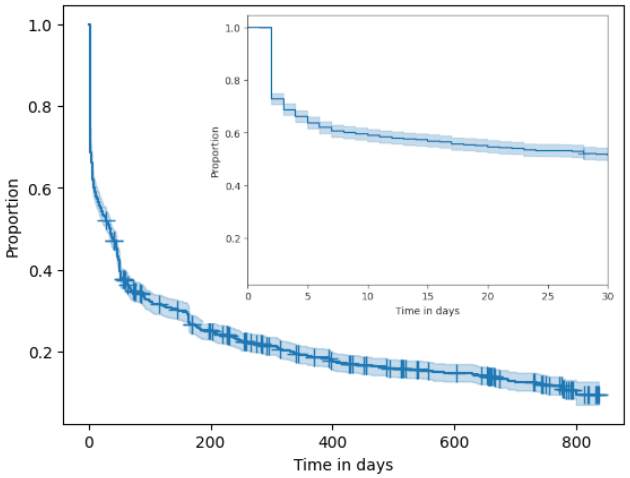
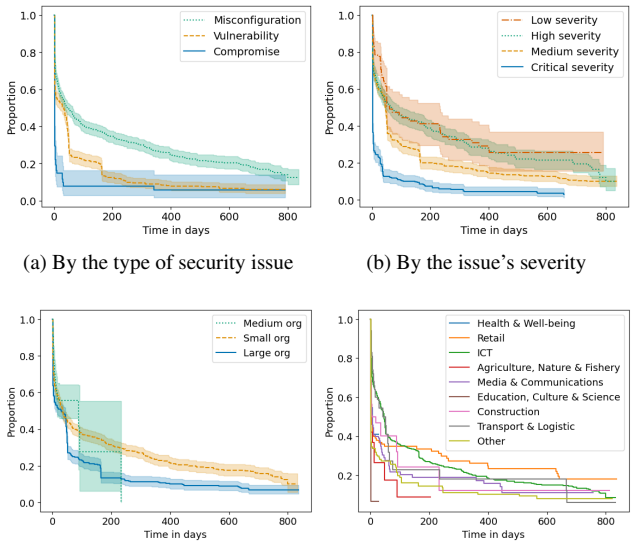


Figure 3: **Kaplan-Meier curve of the remediation estimation of notified security issues**



(a) By the type of security issue

(b) By the issue's severity

(c) By organization's size

(d) By organization's sector

Figure 4: **Kaplan-Meier curve of the remediation estimation based on the different values for the type of security issue, its severity, the organization's sector, and size**

diated. The survivors are the unmitigated cases. Over time, the number of surviving cases goes down. Any case that has not been mitigated after the last sent notification is modeled as censored. Censoring refers to situations where we do not know the exact time a mitigation occurs for a subject, either because the mitigation hasn't happened yet by the end of the study or the 'subject' leaves the study early. We only know that their actual mitigation time is longer than the time we observed. Censoring these cases is necessary because if we were to include all data points as they are and conduct the analysis directly, systems notified near the end of the program would appear unmitigated for the entire duration, even though they had only been active for a few days or months.

From the Kaplan–Meier curve in Figure 3, censored observations (vertical dashes) appear throughout the program, with a concentration toward the end. Approximately 27% of issues are remediated within one day of notification, 40% within one week, and 49% within one month, while about 20% remain affected one year after the initial notification. The estimated median remediation time is 35 days (95% CI: 28–39 days).

At 800 days, the Kaplan–Meier estimator yields a 12% probability of remaining unresolved, corresponding to an 88% remediation probability. This is substantially lower than the 25% raw proportion of unresolved issues at the end of the program. The discrepancy reflects how Kaplan–Meier correctly accounts for censored cases whose remediation may simply occur after monitoring ends, providing a less biased estimate than the raw proportion. In essence, the Kaplan–Meier curve primarily captures the speed and trajectory of remediation rather than just a single end-of-study rate.

In general, we observe higher mitigation rates compared to previous studies, as illustrated in Figure 1. Our measured remediation, even when we exclude notifications about compromise, considerably outperforms those in other vulnerability and misconfiguration notification experiments with the standard setup. The only studies with comparable or better results employed aspects in their setup that are not scalable in the long run, e.g., using postal mail, or relied on statements regarding legal and monetary fines and consequences if the recipient did not mitigate the issue. From this comparison, there are strong reasons to believe that the remediation behavior we observe differs from that of previously studied groups.

These results present the general remediation across the entire dataset; however, they provide no insight into which factors may influence the remediation rate and speed. We begin by examining the differences in notification characteristics. Does the type of issue the notification concerns have any influence on the remediation behavior? We performed a Log-rank test to see if there were significant differences between the three types of notifications for compromise, vulnerability, and misconfiguration. In Figure 4a, we see that organizations seem to prioritize compromises over the other types, with a one-day mitigation rate of 66%. On the other hand, misconfigurations are often mitigated the slowest, with just 20% of issues receiving a fix within a day. This is consistent with what we heard during the interviews.

We also checked the relationship between the severity of the issue, as classified by Shadowserver, and remediation. As we already mentioned, the CSIRT does not include a severity label with their notifications, except for critical severity issues. They leave it up to the recipients of the notifications to interpret and prioritize the other issues. As one might expect, the mitigation behavior for critical security flaws differs

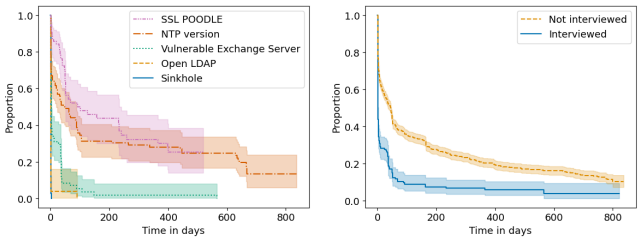
significantly from that for all other severity levels. This is also evident in the Kaplan-Meier plot in Figure 4b, which shows that critical issues are remediated faster, with 56% resolved within a day, compared to around 20% for lower severities. This finding suggests that organizations come to different assessments of severity, especially in the absence of pre-assigned severity labels by the notifier. Another key reason to come to different assessments is that the asset containing the issue might have a different operational value to the organization, leading to the severity being estimated as higher or lower.

When distinguishing the dataset by the organizational characteristics, the results are somewhat mixed. One statistically significant difference is observed in the organizational characteristics between the remediation for small and large organizations, as shown in Figure 4c. Large organizations tend to show significantly faster remediation. This could be attributed to differences in available security resources, knowledge, and personnel. Smaller entities are usually restricted in these aspects. Based on size, the Log-rank test showed that there is no evidence of differences between medium organizations and the other two groups, but small and large organizations do behave significantly differently in terms of remediation, as confirmed by the Kaplan-Meier curves.

Based on the organization's sector, as shown in Figure 4d, most curves overlap. According to the Log-rank test statistics, the only two sectors with a large number of statistically significant differences ($p < 0.05$) compared to other sectors are Education, Culture & Science, and Transport & Logistics. All other sectors show no significant differences. In short, sectoral differences do not seem to be a major driver for mitigation.

As we mentioned before, during the interviews, we asked the participants to prioritize a set of five sample notifications that we presented to them. In that task, we showed them notifications regarding a sinkhole connection (compromise), an accessible LDAP server (misconfiguration), a vulnerable Microsoft Exchange server (vulnerability), a system vulnerable to the SSL POODLE attack (vulnerability), and a misconfigured NTP protocol (misconfiguration). We asked them to rank them in terms of priority and urgency. To compare their answers to what is actually visible in the notification data, we created the remediation plot in Figure 5a. From the plot, it is visible that the issues concerning a sinkhole connection and an accessible LDAP service are usually mitigated the fastest. At the same time, misconfiguration of the NTP protocol or the use of an outdated SSL version has a lower priority. During the interviews, we received various prioritization classifications, some of which align with these findings, most notably that the majority chose to address the compromised asset (the sinkhole notification) as soon as possible.

During our interviews, majority of our respondents reported that they act on all notifications. To compare this claim to our notification data, we analyzed whether the group of organizations we interviewed exhibits a different mitigation pattern



(a) By the notifications from the interviews (b) By the interviewed organizations

Figure 5: **Kaplan-Meier curve estimating remediation rate, based on type and severity of security issue**

compared to the rest of the non-interviewed organizations.

In Figure 5b, as well as from the Log rank test, it is visible that there are significant differences in those two groups. Within a day of the initial notification, the organizations we interviewed remediated 56% of the underlying issues, compared to only 23% for the group of organizations that declined our invitation for a talk. This difference persists even when we look at one week post-notice, with 69% compared to 36%. Our interviewees exhibit faster remediation than the other program participants. This raises two points. First, it provides evidence that our recruitment process for the interviews introduced some selection bias: organizations that performed better in remediation were more likely to accept our invitation to be interviewed. Second, although the group of interviewed organizations claimed to act on all notifications, we do not observe a 100% remediation rate. The actual measured remediation sits at 89%.

What does this mean? Interviewees may have given socially desirable answers and do not really act on all notifications. Another explanation is that the way we measure mitigation involves some degree of measurement error. First, our method does not capture all the mitigation actions. For example, some mitigations do not patch the vulnerability or misconfiguration itself, but instead put in place controls – e.g., firewall rules – to stop the attempt at exploitation. Second, our method includes cases where the recipient organization evaluated that the notification did not concern an asset under their control, e.g., it was for an asset they forgot to de-register with the CSIRT or for an asset of a client, so the recipient cannot act on it themselves. This measurement error might explain why the observed remediation rate peaks at 89%, even though the interviewees say they act on everything. These types of measurement errors plague most, if not all, notification experiments in the scientific literature. In other words, academic measurements likely underestimate the degree of remediation.

We estimated a Cox proportional hazards model to study the association between notification characteristics and remediation time. In our setting, each notification case corresponds to a distinct asset–issue combination, and the number

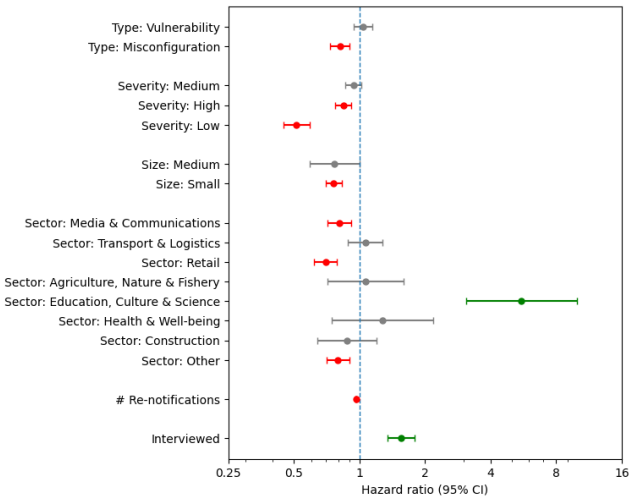


Figure 6: **Cox PH model results compared to the baseline**

of organizations is modest relative to the total number of cases. We therefore treat cases as observational units and fit a conventional Cox model without explicit random effects or cluster-robust variance, following common practice in applied survival analyses when clustering is limited, and the primary interest lies in population-level associations [36]. As in other applications of the Cox model, we acknowledge that some within-organization correlation may persist; however, theoretical and empirical work on robust inference for Cox regression indicates that modest violations of independence primarily affect variance estimates rather than the consistency of coefficient estimates [15]. We therefore adopt the standard Cox model as a pragmatic and widely accepted compromise between model complexity and interpretability in this context. We use the time-varying version of the model due to the time influence of our re-notification covariate, which changes dynamically as time progresses without remediation. Since each re-notification ($\approx$ every 28 days) may alter the likelihood of remediation, the hazard of remediation is not constant over time and depends on this evolving exposure. Using a time-varying Cox model allows the analysis to capture these within-subject changes and accurately estimate the impact of re-notifications on remediation.

The Cox Proportional Hazards regression model uses one value from each covariate as a baseline value. It estimates how the other values of the covariate affect the likelihood of mitigation over time. Hazard ratios indicate whether the likelihood of mitigation is higher or lower compared to the reference value, after controlling for other covariates. In our model, we have used the following baseline values: compromise type, critical severity, large-sized organizations, ICT sector, no re-notifications, and not interviewed. The results of our regression analysis are visualized in Figure 6. Goodness of fit test statistics are in Appendix G.

Our regression model confirms our earlier findings, where we discussed each variable separately, except for one point of nuance regarding the type of issue that the notification is about. The difference between compromise and vulnerability is no longer statistically significant, now that we control for all other covariates, suggesting that other factors in our model might have confounding influences. These findings are consistent with the findings from the interviews that compromise and vulnerabilities that are being exploited in the wild are both prioritized by the interviewed organizations. The difference between compromise and misconfiguration remains significant in the model.

The model confirms all other findings we discussed. Compared to critical severity notifications, all other levels have a reduced likelihood of remediation. Small and medium-sized organizations remediate slower than large organizations. There are no clear differences among the sectors in terms of mitigation. While the sector of Education, Culture & Science shows improvement, this sector contains only a single organization in our sample, so realistically, this effect is not related to the sector. We also observe that re-notifying about an unmitigated issue has a marginal negative effect, compared to the group of notifications where no reminder was sent. In other words, reminder notices are often ignored, especially when the underlying issue has been present for a prolonged period of time. Finally, the model confirms that the group of organizations we interviewed performed better than those who declined our invitation to talk. In sum, coming back to our research question, we have seen an 88% overall remediation rate, with 27% and 49% in shorter time periods of one day and one month, respectively. This rate varies depending on the factors we just discussed.

7 Discussion

Do solicited programs move the needle? As efforts to improve organizational cybersecurity intensify, notifications have become a central intervention deployed by many governments and national CSIRTs. Yet prior work found that unsolicited notices yield limited remediation gains. Our results show that solicited programs can substantially enhance notification effectiveness, making the institutional design of notification mechanisms a critical lever for improving security outcomes.

Organizations that opt in, register their assets, and provide a working contact point, show higher remediation rates compared to prior unsolicited email-based experiments. About 27% of all issues are remediated within a day of notification, about 49% within a month. These levels are well above the 10–30% remediation rates commonly reported for unsolicited vulnerability and misconfiguration notifications over comparable time horizons, and similar to the best-performing – but hard-to-scale – approaches, such as postal mail, instant messaging, or notifications backed by legal threats.

Close to 88% of all issues are eventually resolved. How should we evaluate this number in terms of success? Is there

still room to grow towards 100%? We think this is unlikely. The respondents claimed that they act on all notifications. One might be skeptical and wonder if this claim is driven by a social desirability bias. Yet it is actually quite consistent with our data. The 88% is a lower-bound measurement. The actual remediation rate is likely higher. There are three explanations that suggest the remaining 12% does not reflect a lack of effort.

First, there are non-visible mitigations. Organizations sometimes mitigate risk without changing the externally observable signal—e.g., by adding firewall rules that block exploitation. From the scanner's perspective, the issue persists; from the defender's perspective, the risk may be acceptable or resolved.

Second, our (and the CSIRT's) perspective does not account for errors in ownership and responsibility boundaries. Some assets that no longer belong to the organization, or are operated by clients or service providers. Recipients may correctly decide that "acting" means forwarding or documenting the issue rather than patching the asset themselves. This action is invisible to our measurement.

Third, scan-based detection is imperfect, and errors remain. Vulnerability scans are known to suffer from false positives, especially when they are banner-based [12]. Of course, there are also false negatives. But those errors would not lead to notifications, so they do not impact our results.

These considerations lead us to believe that the 88% is close to the maximum effectiveness one can expect from any notification program. The organizations in our program are arguably a positively selected group: they are security-aware enough to enroll, maintain dedicated contacts, and agree to be interviewed. Solicited programs are the best case of what notification mechanisms can achieve.

The heterogeneity we observe reinforces this point. Compromised systems and critical issues are remediated markedly faster than routine misconfigurations and lower-severity vulnerabilities. Large organizations remediate more quickly than small ones. Interviewed organizations remediate more quickly than those that declined to talk to us. These patterns suggest that solicited programs do not homogenize behavior; instead, they amplify underlying differences in resources, capabilities, and risk perception. The key insight is that notification effectiveness is not a monolithic quantity. It is a distribution that depends on who you notify, about what, and under what institutional arrangement.

7.1 Recommendations

Our interviews and survival analysis point to concrete design choices for more effective, sustainable notification programs.

Target the operational edge, not only top management. Participation in our program was overwhelmingly initiated "from below" by security practitioners, not mandated by executives. Recruitment and communication should therefore focus on operational staff: clear technical value, low administrative overhead, and tooling that fits existing workflows. Regulatory mandates under the EU's NIS2 directive [23] may increase participation, but engagement will hinge on usefulness for those processing the notifications.

Make prioritization support part of the service. Participants triage notifications using perceived exploitability, severity scores, and known campaigns. Compromises and critical issues are remediated quickly; other issues compete for scarce attention. Smaller organizations in particular relied on external sources (e.g., national advisories, CVSS) to gauge urgency. CSIRTs can support this with severity labels.

Support self-management and feedback loops. Participants consistently asked for a portal or dashboard to (*i*) view notification history, (*ii*) maintain asset registrations, and (*iii*) indicate how issues were handled, especially when mitigation is not externally visible or assets are out-of-scope. Lacking such mechanisms frustrates participants and prevents CSIRTs from learning from outcomes and refining targeting. Structured feedback would also reduce measurement blind spots and help distinguish true non-response from artifacts.

Re-think re-notifications and long-lived issues. Our time-varying Cox model shows that re-notifications are at best marginally and sometimes negatively associated with remediation. Once an issue persists over several 28-day cycles, additional reminders are often ignored, suggesting diminishing returns, "alert fatigue," and possibly the absence of a functioning feedback loop, as discussed above. For long-lived issues, CSIRTs may need to change strategy – e.g., escalate to different contacts or de-prioritize these notifications.

Leverage the community dimension. The program also functions as a community platform: participants value the forum and CSIRT days to exchange experiences and report vulnerabilities in third parties. This "reverse flow" turns participants into sensors and amplifiers, extending reach beyond the initial asset list. Investing in community building can materially improve situational awareness and trust.

Our results appear as notification is being institutionalized by law. Under the EU's NIS2 directive, many essential and important entities must maintain contact with governmental or sectoral CSIRTs, with greater emphasis on near real-time notifications of vulnerabilities and compromises. Our findings indicate both opportunities and pitfalls for this regulatory turn.

On the positive side, the solicited program we study resembles a proto-NIS2 world: registered assets, established contact channels, and a trusted national notifier. The improved remediation we observe suggests that institutionalizing such relationships can meaningfully reduce exposure – especially for high-severity and compromise-related issues – without heavy-handed enforcement.

On the negative side, we observe dynamics that regulation alone cannot fix. First, we encountered a false sense of security among some participants. They interpret "no notifications" as evidence of good security. Given the narrow coverage of any single scanning source, this is dangerous. Program de-

sign and regulatory guidance must clarify what notification does not guarantee, to avoid over-reliance on any single feed or service. Second, there are persistent capacity gaps. Small organizations remediate more slowly, even when they opt in and receive the same information. Additional support may be needed. Third, there are asset ownership complexities. Many organizations, especially smaller ones, do not own or directly manage their internet-facing assets. In these cases, service providers – not end customers – should be the primary notification targets. Regulatory frameworks and CSIRT practice must address such multi-layer governance structures if notification is to translate into real risk reduction.

8 Conclusions

In this study, we presented the first-ever empirical measurement of remediation behavior in a solicited security notification program. We found that remediation success is significantly higher than those reported in other efforts. Based on these results and interviews with the program participants, we derived recommendations for improved program design and for regulatory efforts to scale up such programs.

Ethical Considerations

We present this section as a stakeholder-based analysis identifying affected parties, their potential benefits and harms, and the mitigation steps taken to address them. The analysis is built on the four principles of the Menlo Report [1]: Beneficence (do good and minimize harm), Respect for Persons (treat individuals as autonomous agents and protect their ability to make free, informed decisions), Justice (distribute the benefits and burdens of research fairly), and Respect for Law and Public Interest (operate within applicable legal frameworks and serve the broader public good). We apply these principles both to the research process and the implications of our findings post-publication.

Stakeholder identification We identify the following primary and secondary stakeholders: (1) Participating Organizations: These organizations (i.e., enterprises) voluntarily enrolled in the CSIRT's notification program, explicitly requesting to receive security alerts about their own registered assets, which means their security data was already being collected and acted upon before our research began. They are directly affected by the notifications and by how the data is analyzed and reported. (2) IT and Security Personnel within Participating Organizations: Individual practitioners process notifications and decide on remediation actions. Interview participation could expose professional opinions, operational practices, or perceived shortcomings. (3) The Governmental CSIRT: The CSIRT operates the notification program and maintains trust relationships with participants. Research findings could affect perceptions of the program's effectiveness and credibility, as well as future policy decisions. (4) Third Parties and Internet Users: Remediation of vulnerabilities and misconfigurations reduces broader ecosystem risk, benefiting users and other organizations indirectly exposed to compromised or misconfigured systems. This group bears none of the burdens of the research but stands to benefit from improved program design, which is a relevant consideration under the Justice principle. (5) The Research Community and Policymakers: The study informs future notification program design, regulatory approaches (e.g., NIS2), and academic experimentation involving real-world security interventions.

Research impact The publication of this work entails both positive and potential negative impacts.
Positive impact. (1) Improved notification program design: Our findings provide actionable recommendations for CSIRTs and policymakers on key design choices for solicited notification programs, strengthening public–private collaboration and improving recipient value. (2) Bridging communication: By relaying structured feedback from participating organizations to the CSIRT, the research supported concrete program improvements that directly benefit participants. (3) Empirical implications: As the first empirical evaluation of a solicited notification program, this work provides real-world measurement and serves as a reference point for future scientific comparisons of notification effectiveness.
Negative impact. (1) Organizational risk: If improperly handled, findings could expose participating organizations to reputational harm or adversarial targeting. (2) Perceived participation pressure: Collaboration with a governmental CSIRT could create perceived pressure to participate in interviews or provide socially desirable responses, despite voluntary recruitment. (3) Policy misapplication: Policymakers could overgeneralize the observed remediation gains and mandate similar programs without appropriate contextualization or stakeholder engagement.

Mitigation We implemented multiple safeguards to address these risks. (1) All analyzed security issues originate from historical reports routinely produced by the CSIRT and its data providers. The authors did not conduct scanning, discover vulnerabilities, or generate new security findings. No control group was introduced, as withholding notifications would conflict with the CSIRT's legal and ethical obligations and could expose organizations to harm by leaving them unaware of active vulnerabilities. Using pre-existing historical data meant that no participant faced any new risk as a consequence of our research. (2) Quantitative data processing was conducted under a non-disclosure agreement on secure governmental infrastructure with strict access controls. No raw data is shared publicly. All personally identifiable information was removed, and results are reported only in aggregated form to prevent

re-identification or backward inference. This directly mitigates the risk of adversarial targeting, complies with GDPR, and reflects the Justice principle. (3) The interview study received approval from the university's Institutional Review Board (IRB). Participation was voluntary, informed consent was obtained, and participants were reminded of their right to decline to answer questions or to terminate the interview. Prior to publication, participants reviewed quoted material to prevent re-attribution or de-anonymization.

Justification for research This study addresses an important gap in understanding how defensive vulnerability notification programs operate and influence remediation behavior across organizations. Such programs are widely promoted as a mechanism for improving cybersecurity hygiene, yet there is limited systematic evidence on their real-world effectiveness. By analyzing historical program data and conducting voluntary stakeholder interviews with appropriate privacy protections, we provide empirical insights that benefit CSIRTs, enterprises, and the broader community. The benefits of improving defensive practice and informing policy outweigh the mitigated risks of privacy exposure and potential adversarial misuse, particularly given the strict anonymization, aggregation, and ethical safeguards implemented during both data handling and reporting.

Open Science

The authors strongly support the use of open data and the open science initiatives. Yet, access to data used in this paper is restricted due to the sensitive nature of said information and is shared with the authors under a non-disclosure agreement by the governmental CSIRT. As discussed in the 'Ethics' section, restricted and controlled access to the notification data has been granted to the authors on governmental encrypted infrastructure. Therefore, it is not legally possible to make the notification data in this paper public. Nevertheless, the authors enclose the interview protocol and final codebook used for the user study portion of the research in Appendix A and Appendix B, respectively.

References

[1] The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. URL: https://www.dhs.gov/sites/default/files/publications/CSD-MenloPrinciplesCORE-20120803_1.pdf.

[2] Early Warning Service | Arctic Security. URL: https://www.arcticsecurity.com/products/early-warning-service.

[3] ATLAS.ti - Software for Qualitative Data Analysis. URL: https://atlasti.com/.

[4] Virginia Braun and Victoria Clarke. One size fits all? what counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology*, 18(3):328–352, 2021. arXiv:https://doi.org/10.1080/14780887.2020.1769238, doi:10.1080/14780887.2020.1769238.

[5] F.O. Cetin, Carlos Ganán, Maciej Korczynski, and Michel van Eeten. Make notifications great again: learning how to notify in the age of large-scale vulnerability scanning. In *16th Workshop on the Economics of Information Security (WEIS 2017)*, pages 1–23, 2017. Accepted Author Manuscript; WEIS 2017 ; Conference date: 26-06-2018 Through 27-06-2018.

[6] Cyber safety review board report on summer 2023 microsoft online exchange incident | cisa. URL: https://www.cisa.gov/resources-tools/resources/CSRB-Review-Summer-2023-MEO-Intrusion.

[7] Ransomware vulnerability warning pilot (RVWP) | CISA. URL: https://www.cisa.gov/stopransomware/Ransomware-Vulnerability-Warning-Pilot.

[8] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, et al. The matter of heartbleed. In *Proceedings of the 2014 conference on internet measurement conference*, pages 475–488, 2014.

[9] Glossary: Enterprise size - Eurostat. URL: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Glossary%3AEnterprise_size.

[10] FIRST. CSIRT services framework version 2.1. URL: https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1.

[11] Anne Hennig, Fabian Neusser, Aleksandra Alicja Pawelek, Dominik Herrmann, and Peter Mayer. Standing out among the daily spam: How to catch website owners' attention by means of vulnerability notifications. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI EA '22, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3491101.3519847.

[12] Szu-Chun Huang, Harm Griffioen, Max van der Horst, Georgios Smaragdakis, Georgios van Eeten, and Yury Zhauniarovich. Trust but Verify: An Assessment of Vulnerability Tagging Services. In *34th USENIX Security Symposium*, Aug 2025. URL: https://www.usenix.org/conference/usenixsecurity25/presentation/huang-szu-chun.

[13] Frank Li, Zakir Durumeric, Jakub Czyz, Mohammad Karami, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. You've got vulnerability: exploring effective vulnerability notifications. In *Proceedings of the 25th USENIX Conference on Security Symposium*, SEC'16, page 1033–1050, USA, 2016. USENIX Association.

[14] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. Remedying Web Hijacking: Notification Effectiveness and Webmaster Comprehension. In *Proceedings of the 25th International Conference on World Wide Web*, WWW '16, page 1009–1019, Republic and Canton of Geneva, CHE, 2016. International World Wide Web Conferences Steering Committee. doi:10.1145/2872427.2883039.

[15] D. Y. Lin and L. J. Wei. The robust inference for the cox proportional hazards model. *Journal of the American Statistical Association*, 84(408):1074–1078, 1989. URL: https://dlin.web.unc.edu/wp-content/uploads/sites/1568/2013/04/LinWei89b.pdf, doi:10.1080/01621459.1989.10478874.

[16] Qasim Lone, Alisa Frik, Matthew Luckie, Maciej Korczyński, Michel van Eeten, and Carlos Gañán. Deployment of Source Address Validation by Network Operators: A Randomized Control Trial. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 2361–2378, 2022. doi:10.1109/SP46214.2022.9833701.

[17] Max Maaß, Marc-Pascal Clement, and Matthias Hollick. Snail Mail Beats Email Any Day:On Effective Operator Security Notifications in the Internet. In *Proceedings of the 16th International Conference on Availability, Reliability and Security*, ARES '21, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3465481.3465743.

[18] Max Maass, Alina Stöver, Henning Pridöhl, Sebastian Bretthauer, Dominik Herrmann, Matthias Hollick, and Indra Spiecker. Effective Notification Campaigns on the Web: A Matter of Trust, Framing, and Support. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2489–2506. USENIX Association, August 2021. URL: https://www.usenix.org/conference/usenixsecurity21/presentation/maass.

[19] Cybersecurity signals: Connecting Controls and incident outcomes. URL: https://www.marsh.com/en/services/cyber-risk/insights/cybersecurity-signals.html.

[20] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–23, November 2019. URL: https://dl.acm.org/doi/10.1145/3359174, doi:10.1145/3359174.

[21] Early warning program | NCSC. URL: https://www.ncsc.gov.uk/section/active-cyber-defence/early-warning.

[22] Nationaal Cyber Security Centrum Nederland (NCSC-NL). URL: https://www.ncsc.nl/.

[23] Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive). URL: https://eur-lex.europa.eu/eli/dir/2022/2555/oj.

[24] Yevheniya Nosyk, Maciej Korczyński, Carlos H. Gañán, Michał Król, Qasim Lone, and Andrzej Duda. Don't Get Hijacked: Prevalence, Mitigation, and Impact of Non-Secure DNS Dynamic Updates. In *2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 1480–1489, 2023. doi:10.1109/TrustCom60117.2023.00202.

[25] National Vulnerability Database (NVD) - Vulnerability Metrics. URL: https://nvd.nist.gov/vuln-metrics/cvss.

[26] Tara Poteat and Frank Li. Who you gonna call? an empirical evaluation of website security.txt deployment. In *Proceedings of the 21st ACM Internet Measurement Conference*, IMC '21, page 526–532, New York, NY, USA, 2021. Association for Computing Machinery. URL: https://doi-org.tudelft.idm.oclc.org/10.1145/3487552.3487841, doi:10.1145/3487552.3487841.

[27] Takayuki Sasaki, Akira Fujita, Carlos H. Gañán, Michel van Eeten, Katsunari Yoshioka, and Tsutomu Matsumoto. Exposed Infrastructures: Discovery, Attacks and Remediation of Insecure ICS Remote Management Devices. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 2379–2396, 2022. doi:10.1109/SP46214.2022.9833730.

[28] Sezaneh Seymour and Daniel Woods. Calibrating secure by design with the risks faced by small businesses, Feb 2025. URL: https://www.lawfaremedia.org/article/calibrating-secure-by-design-with-the-risks-faced-by-small-businesses.

[29] ShadowServer. URL: https://www.shadowserver.org/.

[30] Shodan monitor. URL: https://monitor.shodan.io/.

[31] Wissem Soussi, Maciej Korczynski, Sourena Maroofi, and Andrzej Duda. Feasibility of Large-Scale Vulnerability Notifications after GDPR. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 532–537, 2020. doi:10.1109/EuroSPW51379.2020.00078.

[32] Ben Stock, Giancarlo Pellegrino, Frank Li, Michael Backes, and Christian Rossow. Didn't You Hear Me? – Towards More Successful Web Vulnerability Notifications. In *Proceedings of the 25th Annual Symposium on Network and Distributed System Security (NDSS '18).*, February 2018. URL: https://publications.cispa.saarland/1190/.

[33] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. Hey, you have a problem: On the feasibility of large-scale web vulnerability notification. In *25th USENIX Security Symposium*, Aug 2016. URL: https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/stock.

[34] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. POSTER: Mapping the Landscape of Large-Scale Vulnerability Notifications. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, page 1787–1789, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2976749.2989057.

[35] Alina Stöver, Nina Gerber, Henning Pridöhl, Max Maass, Sebastian Bretthauer, Indra Döhmann, Matthias Hollick, and Dominik Herrmann. How Website Owners Face Privacy Issues: Thematic Analysis of Responses from a Covert Notification Study Reveals Diverse Circumstances and Challenges. *Proceedings on Privacy Enhancing Technologies*, 2023:251–264, 04 2023. doi:10.56553/popets-2023-0051.

[36] Terry M. Therneau and Patricia M. Grambsch. *Modeling Survival Data: Extending the Cox Model*. Statistics for Biology and Health. Springer, New York, 2000. doi:10.1007/978-1-4757-3294-8.

[37] Christine Utz, Matthias Michels, Martin Degeling, Ninja Marnau, and Ben Stock. Comparing Large-Scale Privacy and Security Notifications. *Proc. Priv. Enhancing Technol.*, 2023(3):173–193, 2023. URL: https://doi.org/10.56553/popets-2023-0076, doi:10.56553/POPETS-2023-0076.

[38] Marie Vasek and Tyler Moore. Do Malware Reports Expedite Cleanup? An Experimental Study. In *5th Workshop on Cyber Security Experimentation and Test (CSET 12)*, Bellevue, WA, August 2012. USENIX Association. URL: https://www.usenix.org/conference/cset12/workshop-program/presentation/Vasek.

[39] 2025 Data Breach Investigations Report. URL: https://www.verizon.com/business/resources/reports/dbir/.

[40] Whisper AI - Professional Voice to Text Transcription. URL: https://whisperai.com/.

[41] Eric Zeng, Frank Li, Emily Stark, Adrienne Porter Felt, and Parisa Tabriz. Fixing HTTPS Misconfigurations at Scale: An Experiment with Security Notifications. In *The 2019 Workshop on the Economics of Information Security*, 2019.

[42] Jia Zhang, Haixin Duan, Wu Liu, and Xingkun Yao. How to Notify a Vulnerability to the Right Person? Case Study: In an ISP Scope. In *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*, pages 1–7, 2017. doi:10.1109/GLOCOM.2017.8253993.

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

[44] Orçun Çetin, Mohammad Hanif Jhaveri, Carlos Gañán, Michel van Eeten, and Tyler Moore. Understanding the role of sender reputation in abuse reporting and cleanup. *Journal of Cybersecurity*, 2(1):83–98, 12 2016. arXiv:https://academic.oup.com/cybersecurity/article-pdf/2/1/83/10833175/tyw005.pdf, doi:10.1093/cybsec/tyw005.

Appendix Overview

Appendix A: Interview Protocol
Appendix B: Code Book
Appendix C: Dataset Variables
Appendix D: Sample Notifications
Appendix E: Program Participants Demographics
Appendix F: Notifications per Sector
Appendix G: Log Likelihood Results of the Survival Model

A Interview Protocol

Opening questions

1. What is your position, and for how long have you been with the organization?

2. How many employees work in the organization?

3. How many employees work at the IT department?

Questions about the notification program

4. Do you know the consideration of the organization and why they applied? Do you know who initiated the sign-up?

5. How often do you get notifications about the same issue from different sources?

6. Do you usually get the information from the CSIRT earlier or later than the other sources?

7. These are the Internet-facing assets the CSIRT knows about your organization. Are they still up to date?

8. How often have you updated these assets with the CSIRT?

Questions about the notification workflow

9. What is the contact point for new notifications?

10. Can you walk me through the step-by-step process of how notifications are processed?

Questions about the notification assessment

11. Given the following 5 notifications, can you assess them like you would when you receive a notification from the CSIRT by ordering them in priority?

12. What actions, if any, would you take for each notification?

Closing questions

13. What are the expectations of the pilot?

14. How have these expectations evolved/changed?

15. Are there improvements to the CSIRT notification system that you'd like to see in the future?

16. Would you like to add anything else?

B Code book

Table 3: **The final code book of the interviews**

Category	Subcategory
Interviewee position	CISO Security Analyst Security Engineer Security Officer System Administrator System Engineer Other
Organization size	Small ($\leq$ 49 FTEs) Medium (50-245 FTEs) Large (250+ FTEs)
ICT department %	None $<$ 20% 20-50% 50+%
Program sign up initiation	Bottom up Top down
Program expectations	External monitoring and notification More security information Other
Program experiences	Few notifications No news means good news Uncertainty due to lack of notifications Participated in other CSIRT initiatives Other
E-mail address for notifications	Security email Front desk email Personal email
Asset status	Up-to-date Small changes Big changes
Update frequency	Often Once Never
CSIRT Notifications	Action step Classification (e.g. issue type, prioritization) Opinions of the notification
Organization's external resources	Security solutions Security websites and forums Vendor mailing lists
Organization's other efforts	Self scanning Regular monitoring Other
Program improvements	Improved communication Stronger community engagement Self-management dashboard or portal Other

C Dataset variables

Table 4: **Key variables in the notification dataset**

Variable	Description
IP Address	Affected IP address in the notification email
Port	Port associated with the affected asset
Security Issue	Type of security issue reported by Shadow Server
Company ID	Unique identifier of the recipient company
First Notification Date	Date of the first notification
All Notification Dates	All notifications for a given asset–issue pair
Organization Size	Self-reported organizational size
Organization Sector	Self-reported sector as defined by the Ministry of Economic Affairs
Issue Severity	Severity level defined by Shadow Server
Re-notifications	Number of re-notifications sent by the CSIRT
Interviewed	Participation in the user study (yes/no)

D Sample notifications

Sample notification: Vulnerable Exchange Server

Dear participant,

From a reliable source we have received information about a potential vulnerable Microsoft Exchange system within your network. The system is not updated with the latest available security updates. Make sure that the indicated system is updated with the latest security updates, and that Extended Protection has been set to 'enabled'.

In Exchange version 2016 and older, Extended Protection is not enabled by default making them vulnerable even with the latest updates installed. Because of this, older versions of Exchange will be marked as 'possible-cve-2024-xxxxx'. This is a limitation of the non-intrusive method of scanning, that is unable to determine if Extended Protection has been enabled or not.

If you have installed the most recent security updates, AND enabled Extended Protection on a relevant version of Exchange, you can ignore this notification. For more information see links 1, 2, 3 below.

1. https://docs.microsoft.com/en-us/exchange/new-features/build-numbers-and-release-dates
2. https://learn.microsoft.com/en-us/exchange/plan-and-deploy/post-installation-tasks/security-best-practices/exchange-extended-protection
3. https://www.shadowserver.org/what-we-do/network-reporting/vulnerable-exchange-server-report

Microsoft Exchange Servers are common targets for malicious actors. Microsoft regularly releases updates for Exchange. We recommend updating your Exchange server with the latest available security updates and if the system has been vulnerable for a critical vulnerability for a longer period of time to verify if any malicious activity has already taken place. If malicious activity was found, we also advise to conduct further research into whether the malicious actors may have established (other) persistent access to systems in your network.

Microsoft releases new updates by default on the second Tuesday of each month for all of their products. If your organization does not yet have a patch management process, or Microsoft products are not included in this process, we recommend to plan a fixed moment after these Tuesdays to patch all Microsoft systems within your network.

For the technical information, see below or the attached email attachment.

==========

"timestamp","ip","port","hostname","tag"
"X","X","X","X","exchange;possible-cve-2024-21410"

==========

If you have any questions in regards to this report, please let us know.

Sample notification: Sinkhole
Dear participant,

From a reliable source we have received a report about a potentially infected system within your network. A connection was detected to a sinkholed server.

This notification contains events (connections) to sinkholes. Sinkholing is a technique whereby a resource (for example a server) used by malicious actors to control malware is taken over and redirected to a benign listener that can (to a varying degree) understand connections coming from infected devices.

Since a sinkhole server is only accessed through previously malicious domain names, only infected systems should be listed in the scan results.

We advise you to check whether the infected system is actually (still) present. If this is the case, it is advisable to immediately check the affected system for the malicious code (malware). It is also advisable to investigate how this malicious code was placed on the system, in order to avoid recurrence.

This information comes from Shadowserver. For more information see: https://www.shadowserver.org/what-we-do/network-reporting/sinkhole-http-events-report/

For the technical information, see below or the attached email attachment.

==========

"timestamp", "ip", "port", "hostname", "protocol", "destination ip", "destinationport", "destination hostname", "infection", "malware family", "tag"
"X","X","X","X","X","X","X","X","adload","adload","X"

==========

If you have any questions in regards to this report, please let us know.

Sample notification: SSL POODLE Attack
Dear participant,

A report has come in from a reliable source about a system on your network with a vulnerability that may allow a malicious party to obtain information from an encrypted connection.

With this so-called POODLE (Padding Oracle On Downgraded Legacy Encryption) attack, a malicious party can use a Man-in-the-Middle (MitM) attack to force a browser to use the SSL 3.0 protocol. This protocol contains a vulnerability that allows the malicious party to decrypt the encrypted traffic over the connection. This way, the attacker can view the data sent over this connection and, for example, capture login credentials.

We recommend disabling SSL 3.0 on the affected system and using at least TLS 1.2.

This information comes from Shadowserver. For more information see: https://www.shadowserver.org/what-we-do/network-reporting/ssl-poodle-report/

For the technical information, see below or the attached email attachment.

==========
"timestamp", "ip", "port", "hostname", "handshake", "cipher_suite", "ssl_poodle", "cert_length", "subject_common_name", "issuer_common_name", "cert_issue_date", "cert_expiration_date"

"X", "X", "X", "X", "X", "TLS_RSA_WITH_RC4_128_SHA", "True", "2048", "X", "X", "X", "X"
==========

If you have any questions in regards to this report, please let us know.

Sample notification: NTP version
Dear participant,

From a reliable source we have received a report about a system with the Network Time Protocol-service (NTP) publicly accessible via port 123 (UDP). NTP is used to synchronize the time for computers in a network.

A malicious actor can abuse publicly accessible NTP-services by executing DoS-attacks on a third system.

We recommend verifying whether this service should be publicly accessible. If this is not the case, please restrict access to the NTP service. This can be done by for example limiting access to only the devices that should need access. It is also recommended to verify if any updates are available for the NTP daemon on the NTP server.

This information comes from Shadowserver. For more information see: https://www.shadowserver.org/what-we-do/network-reporting/ntp-version-report/

For the technical information, see below or the attached email attachment.

==========
"timestamp", "ip", "protocol", "port", "hostname", "version", "system"

"X", "X", "X", "X", "X", "X", "X"
==========

If you have any questions in regards to this report, please let us know.

Sample notification: Publicly Accessible LDAP Server
Dear participant,

On behalf of the CSIRT, we would like to inform you of a serious cybersecurity threat to your company. This is based on current threat data in our possession.

Are you unsure about this message? Visit our official government website to read how you can verify the authenticity of the CSIRT as the sender.
CSIRT website to verify email authenticity

The cybersecurity threat concerns the following:
A system connected to the internet that can be traced back to your (client) organization is (possibly) misconfigured. It involves an LDAP service that is accessible via the public internet and therefore reachable by anyone. As a result, a malicious actor could potentially gain access to sensitive information about your IT infrastructure.

If the internet-connected system is reachable via a UDP port, the vulnerability may allow an amplification attack to be carried out. This means that a malicious actor could potentially misuse your system to perform a Denial of Service (DoS) attack on a third party's system.

From a reliable source, we have been informed that an LDAP service (2) belonging to you or your client is vulnerable to this threat at the following IP address(es) and port(s):

==========

Timestamp: "X", source IP: "X", source ASN: "X", source reverse DNS: "X", transport protocol: "X", source port: "X", LDAP service name: "X"
Timestamp: "X", source IP: "X", source ASN: "X", source reverse DNS: "X", transport protocol: "X", source port: "X", LDAP service name: "X"

==========

What can you (or your client) do?

- Identify which system is running an open LDAP service.
- If possible, shield the service from the internet.
- Implement IP whitelisting.
- Contact your IT service provider if you are unable to take action yourself.

If you have received this notification as an ISP or service provider, it concerns one of your customers. We ask that you inform the customer associated with the mentioned IP address.

We hope this information enables you to resolve the threat. If you were already aware of this issue and have taken action, you may disregard this message. Should you have any questions about this notification, you may reply to this email.

Lastly, we would appreciate it if you could let us know whether this message was useful to you. You can provide feedback preferably through the feedback form available at the link below (1).

(1) *CSIRT feedback website*
(2) *CSIRT website with LDAP information*

Kind regards,
Operator CSIRT

E Program participants demographics

Table 5: **The demographics of all participating organizations**

ID	Size[a]	Sector[b]	ID	Size	Sector	ID	Size	Sector
1*	S	Constr.	20*	L	H&W	39	L	ICT
2*	M	Constr.	21*	L	H&W	40	S	Retail
3	M	Constr.	22	L	H&W	41	L	Retail
4	M	Constr.	23	L	H&W	42*	L	Retail
5*	L	Constr.	24	S	ANF	43	L	Retail
6*	L	Constr.	25*	L	ANF	44	L	Retail
7*	L	Constr.	26*	L	ANF	45	L	Retail
8	L	Constr.	27	L	ANF	46	L	Retail
9	L	Constr.	28	L	ANF	47	L	Retail
10*	M	M&C	29*	L	ECS	48*	L	Other
11	M	M&C	30*	S	ICT	49	L	Other
12*	L	M&C	31*	S	ICT	50*	L	Other
13	L	M&C	32	S	ICT	51	L	Other
14	L	M&C	33	S	ICT	52	L	Other
15	M	T&L	34	S	ICT	53	L	Other
16	M	T&L	35*	M	ICT	54	L	Other
17	L	T&L	36	M	ICT	55	L	Other
18	S	H&W	37	L	ICT	56	L	Other
19	M	H&W	38	L	ICT	57	L	Other

[a] S-Small; M-Medium;L-Large
[b] Self-reported: Constr. - Construction; M&C - Media & Communications;T&L - Transport & Logistics; H&W - Health & Well-being; ANF - Agriculture, Nature & Fishery; ECS - Education, Culture & Science
* Interviewed organizations

F Notifications per sector

Table 6: **An overview of the distribution of notifications between the different category values**

Sector	# Notif.	% Data
Health & Well-being	22	1%
Retail	100	6%
ICT	1,192	68%
Agri., Nature & Fish.	19	1%
Media & Comms.	138	8%
Edu., Cult. & Science	15	1%
Construction	32	2%
Transport & Logistics	100	6%
Other	148	8%

G Log likelihood results of the survival model

Degrees of freedom: 17, **Null distribution:** Chi-squared

	Test statistic	p	-log2(p)
log-likelihood ratio test	777.43	< 0.005	510.02