

## Secondary Publication



Hennig, Anne; Neusser, Fabian; Pawelek, Aleksandra Alicja; u. a.

### Standing out among the daily spam : How to catch website owners' attention by means of vulnerability notifications

Date of secondary publication: 30.06.2026

Version of Record (Published Version), Conferenceobject

Persistent identifier: urn:nbn:de:bvb:473-irb-115844x

#### Primary publication

Hennig, Anne; Neusser, Fabian; Pawelek, Aleksandra Alicja; u. a. (2022): Standing out among the daily spam: How to catch website owners' attention by means of vulnerability notifications, in: Simone Barbosa, Cliff Lampe, Caroline Appert, u. a. (eds.), CHI EA '22: CHI Conference on Human Factors in Computing Systems Extended Abstracts, New York, NY, United States: ACM, S. 1–8, doi: 10.1145/3491101.3519847

#### Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holders.

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by/4.0/legalcode>

# Standing out among the daily spam: How to catch website owners' attention by means of vulnerability notifications

Anne Hennig  
anne.hennig@kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

Fabian Neusser  
fabian.neusser@stud.uni-bamberg.de  
University of Bamberg  
Bamberg, Germany

Alexsandra Alicja Pawelek  
aleksandra.pawelek@student.kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

Dominik Herrmann  
dominik.herrmann@uni-bamberg.de  
University of Bamberg  
Bamberg, Germany

Peter Mayer  
peter.mayer@kit.edu  
Karlsruhe Institute of Technology  
Karlsruhe, Germany

## ABSTRACT

Running a business without having a website is nearly impossible nowadays. Most business owners use content management systems to manage their websites. Yet, those can pose security risks and provide vulnerabilities for manipulations. With vulnerability notifications, website owners are notified about security risks. To identify common themes with respect to vulnerability notifications and provide deeper insight into the motivations of website owners to react to those notifications, we conducted 25 semi-structured interviews. In compliance with previous research, we could confirm that distrust in unexpected notifications is high and, in contrast to previous research, we suggest that verification possibilities are the most important factors to establish trust in notifications. We also endorse the findings that raising awareness for the severity and the complexity of the problems is crucial to increase remediation rates.

## CCS CONCEPTS

• Security and privacy → Social aspects of security and privacy.

## KEYWORDS

web security, CMS vulnerabilities, vulnerability notification, security awareness

### ACM Reference Format:

Anne Hennig, Fabian Neusser, Alexsandra Alicja Pawelek, Dominik Herrmann, and Peter Mayer. 2022. Standing out among the daily spam: How to catch website owners' attention by means of vulnerability notifications. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '22 Extended Abstracts)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3491101.3519847>

## 1 INTRODUCTION

Nowadays, running a business without having a website is nearly impossible because information about goods and services are mainly

retrieved from online resources [5]. Content management systems (CMS) provide default features that make it easy even for laypersons to create and maintain sophisticated websites [4]. According to W3Techs, nearly two thirds of the top 10 Million Websites use a content management system [24]. But CMSs also pose a security risk. Not only can the CMS's frameworks themselves contain vulnerabilities. Also, there is a vast number of plugins and templates that may introduce vulnerabilities [4].

Conțu et al. [4] describe four main threats for open-source content management systems, which are: manipulation of data through SQL injection or parameter manipulation; gaining unauthorized access to confidential data through SQL injections or Cross-Site Scripting (XSS) attacks; phishing; and code execution by exploiting improperly validated input. As a result, websites suffer from different kinds of attacks, like comment spamming, defacement of websites or redirects. Other attacks aim at the search results of the original website. In case of search engine Spam (SEO Spam) or Pharma Hacks, an attacker deploys code on a website to redirect to fake web shops [19, 20]. The manipulation is not visible on the genuine website, but the sites appear in the search engine results as shops selling illegal or banned drugs and medicines, luxurious brand-name clothing, or expensive appliances for cheap. Often, the malicious code is hidden within the CSS files of a website and cannot be easily found – even by skilled developers [19].

Since the problem is not easy to detect and only visible in a website's search results, most website owners have to rely on vulnerability notifications by the security community to be informed about the manipulation. In trying to create suitable vulnerability notifications, with which we could inform the website owners about the security issues, we found a variety of experimental studies that analyzed senders and communication channels as well as content and design of vulnerability notifications for a variety of web-related security issues [1, 3, 7, 12–14, 16, 17, 21–23, 25–30]. Many of these experiments, however, report low remediation rates and problems to reach out to the recipients of their vulnerability notifications. Rather than conducting another experimental study, where website owners are notified about the vulnerability and then remediation rates are measured, we deem it necessary to talk directly to affected website owners and discuss the perception of vulnerability notifications with them.



This work is licensed under a Creative Commons Attribution International 4.0 License.

*CHI '22 Extended Abstracts, April 29-May 5, 2022, New Orleans, LA, USA*  
© 2022 Copyright held by the owner/author(s).  
ACM ISBN 978-1-4503-9156-6/22/04.  
<https://doi.org/10.1145/3491101.3519847>

The motivation of this paper is to answer the following research questions: (1) How did website owners perceive previous web vulnerability notifications? (2) What are suitable senders and communication channels that the website owners deem trustworthy? (3) What aspects should we consider in future notifications to be deemed trustworthy? (4) What – if any – channels do website owners use to actively inform themselves about security incidents?

We used web crawling results to identify German website owners that were affected by a Pharma Hack or a related SEO spam in the past. We then contacted the affected owners and asked for an interview. To our knowledge, none of the experimental studies have conducted *qualitative interviews* with affected website owners, to identify common themes and trust-promoting factors for a vulnerability notification.

By answering our research questions, we will provide deeper and comprehensive insight into the aspects website owners consider when they receive a vulnerability notification. Furthermore, we will suggest suitable (media) channels to raise awareness for the problem in general. With this, our contributions are:

- (1) Presenting the first qualitative study that provides *detailed information* on how website owners perceive vulnerability notifications.
- (2) Providing *comprehensive* insight into website owners' processes to classify a message as trustworthy. Current research results draw a rather heterogeneous picture. With our research we aim at summarizing the key factors for creating trustworthy vulnerability notifications, and identifying less important factors.
- (3) Providing additional trust-promoting factors for successful notification campaigns that were not described in previous literature.
- (4) Identifying media or channels where the security community can distribute awareness materials to gain website owners' attention. Addressing website owners via channels they use is a key factor for a successful awareness campaign. So far no study on vulnerability notifications has included this in their research.

## 2 RELATED WORK

As stated above, factors influencing notification effectiveness have already been subject of experimental studies. Information about website owners' assessments of previous notifications are mainly retrieved from experimental studies [2, 7, 12, 14, 16, 22, 25–27, 29, 30], some of them accompanied by quantitative surveys [1, 13, 17, 21, 23, 28]. In general, qualitative approaches are more applicable for *understanding* opinions, attitudes, or behavior [9]. But only a few studies in the broader context of vulnerability notifications or vulnerability management have conducted qualitative approaches so far. Still, all of them were dedicated to a different target group or to answer different questions: Dietrich et al. [6] interviewed six system operators to understand how they approach security misconfigurations and what they think are the root caused for security misconfigurations; Jenkins et al. [10] used qualitative coding of 356 list emails from the mailing list of the website PatchManagement.org to identify system administrators' behavior towards

patch management; Krombholz et al. [11] conducted expert interviews with seven experienced security auditors to determine which usability challenges hinder people from deploying strong TLS configurations; and Li et al. [15] interviewed 17 system administrators to understand how they manage software updates.

The previous experimental studies showed that some senders or message types worked better than others. However, there was *no statistically significant difference* in the treatment groups' remediation rates, and, in general, remediation rates were still low. Thus, we could conclude that any vulnerability notification works better than non-notifying website owners. But the impact of factors like sender reputation, amount of information, framing, or a suitable notification channel is still unclear, since none of these factors *alone* increase remediation rates significantly [17, 21, 26, 30].

According to Stock et al. [21], three key factors that lead to a successful notification campaign are the email reading rate, awareness rising factors and the aware-to-fix rate. So far, different aspects of these factors have been researched:

*Identifying an appropriate recipient:* WHOIS technical email contact or contact information from the website are preferred by recipients [21]. WHOIS records are often outdated and thus not always recommendable [28]. CERTs could serve as an intermediate recipient but are found ineffective as well [13]. Manually collecting contact information is more effective, but more laborious and also not sufficient in reaching every website owner [17].

*Successfully delivering the notification:* Although sending notifications via letter seems to be more effective [16], sending notifications via email is still the most efficient way for extensive notification campaigns [21].

*Increasing the awareness rising factors:* The reputation of the sender, correct spelling, providing an accurate and detailed description of the problem, providing a clear motivation that is not attached to financial demands, and providing verification possibilities, amongst others, seem to increase trust in notifications [17]. Whereas digitally signing emails through S/MIME does not seem to increase message reads and remediation rates [21]. However, none of the experiments with variations of these factors showed statistically significant differences within the treatment groups.

*Increasing the aware-to-fix rate:* Higher remediation rates were observed when the content of the notification was tailored to the notified population and more details were provided in the message itself [13, 25, 30]. But again, the experiments showed no statistically significant differences within the treatment groups. Some authors also suggest to provide incentives for remediation, like quarantining compromised websites [27, 29], sanctioning website owners [17], or pointing at reputational damage [25, 28], to illustrate the severity of a problem. However, the comparison of different incentives has not been researched, yet.

## 3 METHODOLOGY

This paper aims to better assess website owners' perceptions and opinions, and identify common themes about vulnerability notifications with an inductive approach. We conducted 25 qualitative interviews with German website owners who were affected by SEO spam manipulations. We developed an interview guideline to discuss suitable senders and communication channels as well

as general aspects of notifications and sources for information about security incidents. We also collected socio-demographic data (company size, the interviewee's role in the company and her/his affinity for technology interaction) that were found to influence remediation rates [22, 30].

Between November 2020 and March 2021, our project partner compiled a list of domains in German speaking countries that were affected by a Pharma Hack or a related SEO spam infection. In April 2021, the complete list was sent to associated German project partners or the German State Offices of Criminal Investigations, which were supposed to inform the website owners about the vulnerability.

Between July and September 2021, we contacted 65 German website owners from this list via email, and asked, if we could call them for an interview. We contacted only German website owners, because we supposed that among those the threshold to answer questions from a German university would be lowest. We used the contact information given on their websites. In our request, we introduced ourselves and the research project, and announced that we would call them in the upcoming days. We also provided our email address and phone number so the recipients could opt out or verify the legitimacy of our request. We called the website owners at least three times afterwards. In total, 25 persons agreed to an interview (response rate: 39%).

All interviews were transcribed using verbatim transcription. Any personal data, like names of persons, companies, places, and domain names were anonymized. We used the software MaxQDA to transcribe and code the interviews. To analyze the interviews, we used open coding as described in [2]. A first coder created an initial codebook based on three interviews, which was iteratively improved with a second coder. Even after analyzing 36 percent of our material together, we could not exceed a moderate intercoder agreement ( $\kappa = 0.51$ ), because the length of the coded segments often varied between both coders. As explained in [2], this is a common problem in analyzing semi-structured interview transcripts. We had, however, a high agreement in the application of the codes. Therefore, we coded the remaining 13 interviews independently and checked on each others codes afterwards. We found only a few disagreements, which we could resolve through discussion.

## 4 RESULTS

Although we explicitly sampled for German website owners who we supposed were already notified, by our associated project partners or the State Offices of Criminal Investigations, only 14 interviewees told that they were informed by one of these third parties prior to our interview request. Seven out of 25 discovered the vulnerability on their own, four of the seven were additionally informed by a third party, unrelated to the research project. Table 1 in Appendix A gives an overview of the senders and notification channels of the initial notification. We also included the size of the business, the value for the affinity for technological interaction (ATI) as well as the suitable senders and notification channel, as mentioned by the interviewees, in the table.

As shown in the table, we interviewed mainly website owners of small businesses (15 interviewees), a few medium sized businesses

(four interviewees) and one person who is working for a large company. We could not determine the business size of five interviewees, because the website belonged to a club or was part of a project, where no one was permanently employed.

All interviewees had a medium affinity for technological interaction of at least 3.33, on a scale from 1 (describing low affinity) to 6 (describing high affinity). 17 interviewees had an ATI value between 4.0 and 5.5 and three interviewees were above 5.5 and therefore close to highest affinity values.

### 4.1 What are suitable senders and communication channels that the website owners deem trustworthy?

If asked for a suitable sender, 13 interviewees named at least the institution or authority which informed them initially (see Table 1 in Appendix A). Fourteen interviewees could not think of a suitable sender intuitively or said they did not care. One interviewee expressed his general mistrust in any sender or channel: "To be honest, I do not know what to do. I mean, there are police officers calling and telling you, they are police officers, and then they scam grandmas [...]." Another told that since emails are often faked, the subject and the content of the message are more important than the sender. Two others confirmed and added that a notification by any sender is better than none.

Eleven interviewees described rather general sender requirements: the sender should be trustworthy, familiar or a government agency, and the email address should be verifiable or fit the sender. Only a few interviewees named specific senders like the Federal Office for Information Security, the Federal Ministry for Economic Affairs and Climate Action, the Federal Ministry of Health, or research institutes like the Helmholtz Association or the Fraunhofer Society. One interviewee said Joomla, the provider of her/his CMS, is a suitable sender, and another interviewee would deem their web agency suitable. Another interviewee also described Karlsruhe Institute of Technology (KIT) as a trust-promoting sender, because her/his son had studied there. One interviewee said s/he did not trust the initial notification because s/he could not see a relation between a private company writing on behalf of a research project. All in all, we noticed that if the participant named a specific sender, s/he had some connection to this sender (e. g., the participant naming the Federal Ministry of Health is working in the health sector).

The same holds for suitable communication channels: 16 interviewees named at least the initial notification channel as suitable (see Table 1 in Appendix A). Notifications via email were deemed suitable by 17 interviewees. Mainly because they value them as easy, fast, and straightforward. Furthermore, emails can contain a lot of information and, as one interviewee explained, "[a]n email always arrives and then, when I see it, I can work through it." One participant even said s/he would have rather received an email than a phone call. Two other interviewees said that email would be the most logical notification channel because it would have been odd to receive information about a digital problem via analog notification channels (one named fax or letter as unsuitable). Five interviewees said they would be more suspicious of email notifications, but no one explicitly refused to be notified via email.

The second most named channel is phone (11 interviewees). Interestingly, most interviewees said they find it easier to verify the trustworthiness of a sender via phone. When phone calls were found unsuitable, it was mainly because the interviewees deemed them too laborious. Only one participant said s/he would not take a notification via a phone call seriously at all. Two interviewees said they think a combination of phone calls (to install trust in the sender) and emails (for further information) would be best.

Only one interviewee said s/he would only trust a notification via letter, and another interviewee thought a letter would be the most genuine notification channel. Notifications via messenger or social media were deemed unsuitable by three interviewees, or were not mentioned at all.

#### 4.2 How were previous notifications perceived, and what should a future notification look like to be considered trustworthy?

Nearly all of the interviewees recognized the initial notification as spam or told us that they mistrust those notifications in general. Five interviewees explicitly said that they classified the notification as spam, while others described the notification as irritating or odd, or they expressed surprise or doubts. Seven interviewees criticized that the information was not sufficient or not comprehensible; therefore, they categorized the notification as spam. Further issues were a suspicious attachment (Excel sheet), a non-matching area code, or the police calling from a cell phone number. Another one said that the email did not look legitimate because it had no signature.

Only one interviewee described the notifications as helpful and said that although the sender did not name the cause for the problem, the notification itself caught her/his attention. One other interviewee was notified by her/his hosting provider. S/he deemed the notification trustworthy and forwarded it to the external service provider. But they declared it a false alert and the company did nothing to close the vulnerability. In another case, the interviewee detected the problem her-/himself, but the hosting provider could not reproduce it and therefore declared it a false alert as well.

We asked the interviewees which aspects of an email notification they deem crucial, and also incorporated factors, which convinced the interviewees to act on the initial notification. We categorized those answers into three categories:

**General Factors:** For the initial notification, **verification possibilities** played a major role in convincing recipients to finally deem the notification trustworthy and act on it. Except for three, all interviewees described some approaches to verify notifications. Seven interviewees verified the problem by reconstructing it. Nine interviewees verified the sender of the notifications by googling, and eight interviewees called the sender, or answered her/him via email. Five interviewees also verified technical details like the sender's domain address. Three interviewees emphasized that some possibility for verification should be included in future notifications.

Ten interviewees expressed the unspecific demands that the **sender** must be reputable, professional, trustworthy, and well-known. The sender's address must also fit the contact information and should look legitimate.

One interviewee said that s/he would prefer a local sender. Another said it was helpful that the police called her/him with a local phone number. On the other hand, no other interviewee explicitly named **geographical proximity** as a trust-promoting factor, and one interviewee even denied that proximity would have made any difference for her/him.

Six interviewees said the notification, in general, should be **plausible**. One interviewee mentioned that the message would be trustworthy if the *hosting provider* sends a message *about defect scripts on the website* and includes the *customer ID* in the subject. Another interviewee said that general plausibility is even more important than a particular sender.

Some interviewees said that it was helpful to be notified by **several entities** or **several times**. One interviewee was notified by our project partner and the police. S/he explained that the police *also* calling her/him made the problem more prominent. Another participant explained that a second email with a more detailed description induced her/him to react and illustrated her/him the seriousness of the problem. Furthermore, two interviewees said that a combination of a phone call and an email would be most suitable for future notifications.

**Content-Related Factors:** A proper **problem description** was another factor for some interviewees to take the initial notification seriously. Another interviewee said that s/he had a "very profound and reliable" conversation with a policeman, which was trust-promoting. Another interviewee emphasized several times that the policeman who notified her/him was very knowledgeable and able to explain the problem adequately. A comprehensive, specific, and informative description of the problem was also the most named content-related factor for future notifications. One interviewee explained that *website abuse* "is such a huge term," and s/he appreciated a detailed instruction. In total, 12 of our interviewees found it trust-promoting that the problem was described in a way that they could understand the issue. Two more interviewees would also like to have information on **how to solve the problem**. Furthermore, one interviewee said that the problem's urgency should be made clear.

One interviewee said, a screenshot that **visualizes the problem** prompted her/him to take action on the initial notification: "I got your email a week ago. But that was ... not so specific [...]. But then you addressed me personally and added the screenshot about the security breach. This really motivated me to take action quickly."

Eight interviewees further asked for providing a clear **motivation**: "Which jurisdictions does the authority have, why does she contact me, what is the objective of the notification?" Regarding the initial notification, one interviewee said that s/he could make sense of the information because the sender mentioned the keyword *cybercrime* right at the beginning of the conversation. Another interviewee said that s/he finally trusted the sender because their motivation was to merely inform her/him about the incident and that s/he would not face criminal charges. One participant said that non-monetary interests *in combination* with the individual salutation and the sender's reputation were trust-promoting for her/him. S/he said: "a personal email, research assistant, research project, KIT. There it all adds up that this is a reliable sender who is not primarily interested in earning money [...]."

On the other hand, the information should **not be too extensive**. One interviewee said that the information should be brief, concise, and tailored to her/his situation, because “messages that are too long are also not read.”

Further content-related factors besides the problem description were a **personalized salutation** (four interviewees), **correct orthography**, and a **meaningful subject** (each named by two interviewees). Two interviewees further said that a customer ID in the subject would be helpful, for example in a vulnerability notification from their webhoster.

*Technical Factors:* Eight interviewees asked for **contact information** in terms of a phone number or email address, a signature, a letterhead, or an imprint. Six interviewees asked for a reasonable, trustworthy, or relatable **domain** in the sender’s email address. One participant said a **link to a website with more information** would be helpful, and another remarked that links are only trust-promoting when the link text matches the link URL.

On the other hand, three interviewees said they never click on links or they would classify notifications with links immediately as spam. One participant named a **digital signature** as trust-promoting but would not require encryption. Two other participants asked for **technical possibilities to verify the authenticity** but did not specify their demands.

### 4.3 What channels do website owners use to inform themselves about security incidents?

Only one interviewee said that s/he subscribed to a mailing list from a publishing house with an IT focus to keep informed about security incidents. Eight interviewees named unspecific channels, like the media in general, on-topic web blogs or web research, and/or the personal environment (business partners, external service provider, colleagues, and friends) as information channels for IT security topics. Four interviewees had installed a malware scanner on their websites. One interviewee said that s/he also uses the website of the Federal Office for Information Security occasionally to find information about specific security incidents. One interviewee said that s/he would watch out for that kind of information in the newsletter of her/his branch association.

## 5 DISCUSSION

In our interviews, we could confirm a high distrust in vulnerability notifications. We could also show that recipients need to connect to the sender or the notification. The decision whether a sender is trustworthy, is thus pretty distinct. Only a few interviewees could think of suitable senders. Some interviewees said sender is irrelevant, since emails and even phone calls can be manipulated. Other interviewees also said, the information itself is most important.

We could not detect any obvious correlations between business size, and the interviewees responses. We might suggest that interviewees with high ATI values, and therefore a high affinity for technological interaction, are more indifferent in their assessment of possible senders. But since our sample is too small and too diverse for statistical analysis, and the focus of this papers is on qualitative interviews, these findings need to be proven by future research.

Maass et al. [17] concluded that no single factor consistently increases trust. We can confirm these findings and can, furthermore, show that even if a sender itself (like the police) or a notification channel (like email) is deemed suitable, the notification process **as a whole** must be plausible to establish trust in the notification. Since previous research could not clearly identify an effective sender and/or notification channel, we conclude that the whole notification process, composed of the sender **and** the notification channel **and** the content of the message must be **reasonable and verifiable**.

*Verification possibilities are most important.* Our main finding is that the recipients seek some way to verify the notification. Thus, we deem it necessary to include some possibilities to verify the information in a vulnerability notification by providing contact information and/or technical possibilities to verify the sender’s authenticity – or by giving the recipients the possibility to reproduce the incident. A clear description of the problem, a plausible motivation for the notification, and, if applicable, information to solve the problem are deemed trust-promoting.

Regarding notification channels, we could show that notifications via email were accepted by most of the interviewees. One participant explicitly stated that email is the most plausible notification channel for a digital process. We, therefore, align with the conclusions of [21] and conclude that sending notifications via email is a justifiable notification channel for large-scale notification campaigns. Like [16], we also suggest considering more laborious notification channels like letters or phone calls for more severe problems, or to *emphasize* severity. It has been shown, that remediation rate are higher if the problem is deemed severe [7, 21, 26].

*Awareness for the problem is important.* As stated in [17, 21, 28], we also could observe that trust in and awareness of a notification does not always cause remediation. As mentioned in Section 4, two interviewees did not see the severity and therefore did nothing to solve the problem, although they deemed the initial notification trustworthy. We, therefore, endorse the suggestions of [16, 17, 25, 28] and highly recommend providing some incentive for remediation or name potential negative consequences from inaction in the notification since some interviewees underestimated the severity of the problem. We further found that either the police as the sender, illustrative materials like screenshots, or demonstrating the consequences can be helpful. One interviewee also mentioned that it could be helpful to have a dedicated location, where information about the problem and counter-measures are given.

Something we could identify in our interviews and in previous research [3] is that the problem is very complex, especially - but not exclusively - for users with little IT (security) background. One interviewee said, s/he talked to her/his hosting provider and they could not reproduce the problem. So s/he dismissed the notification. We, therefore, suggest that it is also necessary to explain how attackers intrude into systems and what mechanisms they use to hide. Furthermore, it is necessary to explain that just because the redirection in the search engine results is temporarily inactive the vulnerability is not automatically fixed and can be misused again.

We could further see that some companies cooperate with external IT providers or agencies. Only four of the 25 interviewees identified their role in the company as system administrators or

employ an administrator in their company. Eight interviewees created the websites themselves or for a relative and 13 participants said they assigned the problem to an external service provider. We, therefore, suggest that it is not sufficient to raise awareness among the affected business owners. It is also necessary to raise awareness amongst hosting providers, as previously suggested in [2, 27]. Furthermore, it is crucial to raise awareness for the severity and the complexity of the problem amongst external IT service providers and web agencies so that they can advise their customers.

External service providers, friends, and colleagues are also channels to exchange information about security incidents. Word-of-mouth, therefore, seems to be a fruitful channel for distributing information to raise awareness. Most of the interviewees we asked did not actively use or seek information about security vulnerabilities. It is, therefore, necessary to also reach out to channels the website owners use unconsciously, like news media or industry associations, to raise awareness for the existence, the severity, and the complexity of the problem.

Due to the qualitative nature of our study, the general validity of our results is limited. So far, we mainly reached out to small and medium-sized businesses in Germany. We also assume that our sample is subject to self-selection bias because we only talked to website owners who we could reach by phone and who agreed to an interview. Nevertheless, we are certain that our findings can be adapted to other contexts, since they are in line with previous quantitative research and – on top of that – provide an explanation why the efficacy of previous notification campaigns is limited.

## 6 CONCLUSION

Previous quantitative research found, that the sender of a vulnerability notification and its reputation seems to play an important role (i.a. [17, 21, 26, 30]). But still, the impact of factors like sender, sender reputation or other factors is not entirely clear, since none of these factors was able to increase remediation rates significantly.

In 25 qualitative interviews with affected website owners, we were able to identify common themes concerning vulnerability notifications. With our work, we could verify existing research and, by using semi-structured interviews instead of quantitative surveys, summarize key factors for creating trustworthy vulnerability notifications, and identify less important factors.

In contrast to previous research, we suggest that sender reputation and/or base trust in a notification channel are important, but not the most important factors in establishing trust in notifications. Most of our interviewees could not name a sender and/or channel they would trust without further request. But nearly all interviewees took some steps to verify either the problem, the sender or the message's content.

As researched by [17], formal and content-related aspects of a notification increase its perceived trustworthiness. We looked at these factors in more detail and could show that especially a clear description of the problem, a clear motivation for the notification, and, if applicable, information to solve the problem should be included in a notification. These factors enable the recipients to verify the problem. Providing contact information (a phone number or email address, a signature, a letterhead, or an imprint) and using

a well-known domain in the sender's email address helps recipients verify the notification's sender. Future notifications should also consider a personalized salutation, correct orthography and a meaningful subject. These factors help recipients to establish a connection to the sender, which, again, helps them to verify the information.

We also agree with [17, 18, 25, 28] and suggest providing some incentive for remediation, like taking a website down temporarily [27, 29], or naming potential negative consequences from inaction, like prosecution [17] or reputational damage [25, 28], in future notification. We also suggest raising awareness for the severity and the complexity of the problems among external IT service providers and hosting providers as well as in the news media in general. These are the channels where affected owners inform themselves about security incidents and where they seek help when they receive a notification or detect a problem themselves.

## ACKNOWLEDGMENTS

This research is supported by the Federal Ministry of Education and Research, Germany, as part of the BMBF INSPECTION project (Zuwendungsnummer 16KIS1113), and it is supported by funding from the topic Engineering Secure Systems, topic 46.23.01 Methods for Engineering Secure Systems, of the Helmholtz Association (HGF) and by KASTEL Security Research Labs.

## REFERENCES

- [1] Jan M. Ahrend, Marina Jirotko, and Kevin Jones. 2016. On the Collaborative Practices of Cyber Threat Intelligence Analysts to Develop and Utilize Tacit Threat and Defence Knowledge on Existing Practices, Shortcomings, System Circumventions and Implications for Design. In *2016 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (CyberSA)*. 1–10. <https://doi.org/10.1109/cybersa.2016.7503279>
- [2] John L. Campbell, Charles Quincy, Jordan Osserman, and Ove K. Pedersen. 2013. Coding In-depth Semistructured Interviews: Problems of Unitization and Inter-coder Reliability and Agreement. *Sociological Methods & Research* 42, 3 (2013), 294–320. <https://doi.org/10.1177/0049124113500475>
- [3] Davide Canali, Davide Balzarotti, and Aurélien Francillon. 2013. The role of web hosting providers in detecting compromised websites. In *Proceedings of the 22nd International Conference on World Wide Web*. Association for Computing Machinery, New York, NY, USA, 177–188. <https://doi.org/10.1145/2488388.2488405>
- [4] Cosmin A. Conțu, Eduard C. Popovici, Octavian Fratu, and Mădălina G. Berceanu. 2016. Security issues in most popular content management systems. In *2016 International Conference on Communications (COMM)*. 277–280. <https://doi.org/10.1109/iccomm.2016.7528327>
- [5] Statista Research Department. 2021. Anteil der Personen, die das Internet zur Suche nach Informationen über Waren und Dienstleistungen genutzt haben, in den Ländern der Europäischen Union (EU-28) im Jahr 2020. <https://de.statista.com/statistik/daten/studie/806662/umfrage/internetsuchenach-informationen-ueber-waren-und-dienstleistungen-in-der-eu/>
- [6] Constanze Dietrich, Katharina Krombholz, Kevin Borgolte, and Tobias Fiebig. 2018. Investigating System Operators' Perspective on Security Misconfigurations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. Association for Computing Machinery, New York, NY, USA, 1272–1289. <https://doi.org/10.1145/3243734.3243794>
- [7] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, and J Alex Halderman. 2014. The Matter of Heartbleed. In *Proceedings of the 2014 Conference on Internet Measurement Conference, IMC '14 (IMC '14)*. Association for Computing Machinery, Vancouver, BC, Canada, 475–488. <https://doi.org/10.1145/2663716.2663755>
- [8] Thomas Franke, Christiane Attig, and Daniel Wessel. 2018. A Personal Resource for Technology Interaction: Development and Validation of the Affinity for Technology Interaction (ATI) Scale. *International Journal of Human-Computer Interaction* 35, 6 (2018), 456–467. <https://doi.org/10.1080/10447318.2018.1456150>
- [9] Cornelia Helfferich. 2011. *Die Qualität qualitativer Daten, Manual für die Durchführung qualitativer Interviews*. VS Verlag für Sozialwissenschaften, Wiesbaden. 1 – 211 pages. <https://doi.org/10.1007/978-3-531-92076-4>

- [10] Adam Jenkins, Pieris Kalligeros, Kami Vaniea, and Maria K. Wolters. 2020. "Anyone Else Seeing this Error?": Community, System Administrators, and Patch Information. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*. 105–119. <https://doi.org/10.1109/EuroSP48549.2020.00015>
- [11] Katharina Krombholz, Wilfried Mayer, Martin Schmiedecker, and Edgar Weippl. 2017. "I Have No Idea What I'm Doing" - On the Usability of Deploying HTTPS. In *26th USENIX Security Symposium (USENIX Security 17)*. USENIX Association, Vancouver, BC, 1339–1356. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/krombholz>
- [12] Marc Kührer, Thomas Hupperich, Christian Rossow, and Thorsten Holz. 2014. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *23rd USENIX Security Symposium (USENIX Security 14)*. USENIX Association, San Diego, CA, 111–125. <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/kuhrer>
- [13] Frank Li, Zakir Durumeric, Jakub Czyw, Mohammad Karami, Michael Bailey, Damon McCoy, Stefan Savage, and Vern Paxson. 2016. You've Got Vulnerability: Exploring Effective Vulnerability Notifications. In *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association, Austin, TX, 1033–1050. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/li>
- [14] Frank Li, Grant Ho, Eric Kuan, Yuan Niu, Lucas Ballard, Kurt Thomas, Elie Bursztein, and Vern Paxson. 2016. Remediating Web Hijacking: Notification Effectiveness and Webmaster Comprehension. In *Proceedings of the 25th International Conference on World Wide Web (WWW '16)*. International World Wide Web Conferences Steering Committee, Montreal, Quebec, Canada, 1009–1019. <https://doi.org/10.1145/2872427.2883039>
- [15] Frank Li, Lisa Rogers, Arunesh Mathur, Nathan Malkin, and Marshini Chetty. 2019. Keepers of the Machines: Examining How System Administrators Manage Software Updates. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. USENIX Association, Santa Clara, CA, 273–288. <https://www.usenix.org/conference/soups2019/presentation/li>
- [16] Max Maass, Marc-Pascal Clement, and Matthias Hollick. 2021. Snail Mail Beats Email Any Day: On Effective Operator Security Notifications in the Internet. In *The 16th International Conference on Availability, Reliability and Security (ARES 2021)*. ACM, New York, NY, USA, Vienna, Austria, 1–13.
- [17] Max Maass, Alina Stöver, Henning Pridöhl, Sebastian Bretthauer, Dominik Herrmann, Matthias Hollick, and Indra Spiecker. 2021. Effective notification campaigns on the web: A matter of Trust, Framing, and Support. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 2489–2506. <https://www.usenix.org/conference/usenixsecurity21/presentation/maass>
- [18] Max Maaß, Henning Pridöhl, Dominik Herrmann, and Matthias Hollick. 2021. Best Practices for Notification Studies for Security and Privacy Issues on the Internet. In *The 16th International Conference on Availability, Reliability and Security (The 16th International Conference on Availability, Reliability and Security)*. Association for Computing Machinery, Vienna, Austria, 1–10. <https://doi.org/10.1145/3465481.3470081>
- [19] Malcare. 2021. What is WordPress Pharma Hack & How to clean it? <https://www.malcare.com/blog/what-is-pharma-hack-how-to-clean-it/>
- [20] Art Martori. 2020. Spamdexing: What is SEO Spam and How to Remove It. <https://blog.sucuri.net/2020/02/spamdexing-seo-spam.html>
- [21] Ben Stock, Giancarlo Pellegrino, Frank Li, Michael Backes, and Christian Rossow. 2018. Didn't You Hear Me? - Towards More Successful Web Vulnerability Notifications. In *Proceedings of the 25th Annual Symposium on Network and Distributed System Security (NDSS '18)*. <https://publications.cispa.saarland/1190/>
- [22] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. 2016. Hey, You Have a Problem: On the Feasibility of Large-Scale Web Vulnerability Notification. In *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association, Austin, TX, 1015–1032. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/stock>
- [23] StopBadware and Commtouch. 2012. Compromised Websites: An Owner's Perspective. <https://www.stopbadware.org/files/compromised-websites-owners-perspective.pdf>
- [24] W3Techs Web Technology Surveys. Usage statistics of content management systems. [https://w3techs.com/technologies/overview/content\\_management](https://w3techs.com/technologies/overview/content_management)
- [25] Marie Vasek and Tyler Moore. 2012. Do Malware Reports Expedite Cleanup? An Experimental Study. In *5th Workshop on Cyber Security Experimentation and Test (CSET '12)*. USENIX Association, Bellevue, WA, 1 – 8. <https://www.usenix.org/conference/cset12/workshop-program/presentation/vasek>
- [26] Eric Zeng, Frank Li, Emily Stark, Adrienne Porter Felt, and Parisa Tabriz. 2019. Fixing HTTPS Misconfigurations at Scale: An Experiment with Security Notifications. In *The 2019 Workshop on the Economics of Information Security (2019)*. Boston, MA, 1 – 19.
- [27] Orçun Çetin, Lisette Altena, Carlos Gañán, and Michel van Eeten. 2018. Let Me Out! Evaluating the Effectiveness of Quarantining Compromised Users in Walled Gardens. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 251–263. <https://www.usenix.org/conference/soups2018/presentation/cetin>
- [28] Orçun Çetin, Carlos Hernandez Ganan, Maciej Korczynski, and Michel van Eeten. 2017. 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)*. WEIS 2017, San Diego, 1–23.
- [29] Orçun Çetin, Carlos Gañán, Lisette Altena, Samaneh Tajalizadehkhoob, and Michel van Eeten. 2019. Tell Me You Fixed It: Evaluating Vulnerability Notifications via Quarantine Network. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. 326–339. <https://doi.org/10.1109/eurosp.2019.00032>
- [30] Orçun Çetin, Mohammad Hanif Jhaveri, Carlos Gañán, Michel van Eeten, and Tyler Moore. 2016. Understanding the role of sender reputation in abuse reporting and cleanup. *Journal of Cybersecurity* 2, 1 (2016), 83–98. <https://doi.org/10.1093/cybsec/tyw005>

## A APPENDIX

Table 1: Overview of sender and notification channels of the initial notification

Interview No	Number of Employees	ATI	Initial Sender	Suitable Sender	Initial Notification Channel	Suitable Notification Channel
1	- <sup>1</sup>	4.22 2	authors	BSI <sup>3</sup> , relatives with IT-Know-How, police, research facilities	e-mail	letter
2	10-49	4.33	project partner	don't know, well-known sender, BMWI <sup>4</sup> , BSI, data protection authorities, IT security, business association, police, hosting provider, research facilities	phone call	letter, e-mail, web portal <sup>5</sup>
3	1-9	3.33	police, project partner	don't know, police, IT security companies	e-mail, phone call	letter, e-mail
4	1-9	4	relatives, project partner	indifferent, well-known and verifiable e-mail address	e-mail, phone call	suspicious of every channel, e-mail
5	1-9	5	self	indifferent, hosting provider, police	-	e-mail
6	-	5.11	project partner, police	BSI	e-mail, phone call	e-mail, phone call
7	1-9	3.44	hosting provider	hosting provider	e-mail	phone call
8	1-9	4	unknown	authorities in general, accounting firm	e-mail	phone call + e-mail
9	1-9	5.33	self	hosting provider, Joomla	-	e-mail
10	1-9	4.78	police	don't know, police, hosting provider	e-mail	letter
11	1-9	4.22	police	indifferent	phone call	phone call
12	-	3.22	authors	content more important than sender, some official body, hosting provider, police, research facilities	e-mail	letter, phone call
13	10-49	4.22	website users	our agency	unknown	phone call
14	1-9	5.44	uninvolved third party	indifferent, police	e-mail	e-mail, phone call
15	1-9	4.44	police	police	phone call and e-mail	e-mail
16	1-9	4.44	hosting provider	hosting provider would be fine	e-mail	e-mail, web portal <sup>6</sup>
17	-	5.78	authors	indifferent, research facilities, hosting provider	e-mail	e-mail
18	1-9	4	authors	don't know, authorities, business association, BMG <sup>7</sup> , city council, police, research facilities	e-mail	e-mail
19	1-9	4.56	self, website users	well-known sender, research facilities	unknown	e-mail
20	-	4.44	police	indifferent, has to match the signature	phone call	e-mail, phone call
21	10-49	5.67	police	don't know, BSI, police	e-mail	letter, e-mail
22	1-9	3.78	police	police, official organizations	phone call	phone call + e-mail
23	1-9	5.67	self	official organizations, hosting provider	-	phone call
24	10-49	4.89	police	don't know, official organizations, police	phone call	don't know
25	50-249	3.33	police	don't know, police	phone call	phone call

<sup>1</sup>If it was not a company website, we could not determine company size by the number of employees<sup>2</sup>Determined via the ATI scale [8] with a 6-point Likert scale from 1.0 to 6.0, with 1.0 meaning low interest in technology and 6.0 meaning a high interest in technology<sup>3</sup>Federal Office for Information Security<sup>4</sup>Federal Ministry for Economic Affairs and Climate Action<sup>5</sup>Only for vulnerability notifications from hosting provider<sup>6</sup>only hosting provider<sup>7</sup>Federal Ministry of Health