From Chaos to Consistency: The Role of CSAF in Streamlining Security Advisories

Julia Wunder
IT Security Infrastructures Lab,
Friedrich-Alexander Universität
Erlangen-Nürnberg (FAU)
Erlangen, Germany
julia.wunder@fau.de

Janik Aurich
IT Security Infrastructures Lab,
Friedrich-Alexander Universität
Erlangen-Nürnberg (FAU)
Erlangen, Germany
janik.aurich@fau.de

Zinaida Benenson
IT Security Infrastructures Lab,
Friedrich-Alexander Universität
Erlangen-Nürnberg (FAU)
Erlangen, Germany
zinaida.benenson@fau.de

ABSTRACT

Security advisories have become an important part of vulnerability management. They can be used to gather and distribute valuable information about vulnerabilities. Although there is a predefined broad format for advisories, it is not really standardized. As a result, their content and form vary greatly depending on the vendor. Thus, it is cumbersome and resource-intensive for security analysts to extract the relevant information. The Common Security Advisory Format (CSAF) aims to bring security advisories into a standardized format which is intended to solve existing problems and to enable automated processing of the advisories. However, a new standard only makes sense if it can benefit users. Hence the questions arise: Do security advisories cause issues in their current state? Which of these issues is CSAF able to resolve? What is the current state of automation?

To investigate these questions, we interviewed three security experts, and then conducted an online survey with 197 participants. The results show that problems exist and can often be traced back to confusing and inconsistent structures and formats. CSAF attempts to solve precisely these problems. However, our results show that CSAF is currently rarely used. Although users perceive automation as necessary to improve the processing of security advisories, many are at the same time skeptical. One of the main reasons is that systems are not yet designed for automation and a migration would require vast amounts of resources.

CCS CONCEPTS

• **Security and privacy** → **Vulnerability management**; *Usability in security and privacy*; • **General and reference** → *Empirical studies.*

KEYWORDS

CSAF, Common Security Advisory Format, Security Advisories, IT Security, Survey, User Study

EuroUSEC 2024, September 30-October 1, 2024, Karlstad, Sweden

ACM ISBN 979-8-4007-1796-3/24/09
https://doi.org/10.1145/3688459.3688463

ACM Reference Format:
Julia Wunder, Janik Aurich, and Zinaida Benenson. 2024. From Chaos to Consistency: The Role of CSAF in Streamlining Security Advisories. In *The 2024 European Symposium on Usable Security (EuroUSEC 2024), September 30-October 1, 2024, Karlstad, Sweden.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3688459.3688463

1 INTRODUCTION

The IT sector is growing steadily from year to year, and the number of connected devices is increasing. By 2030, it is estimated that more than 29 billion devices will be connected [13]. However, the complexity increases as well, as many systems build and depend upon each other. No system is perfect and thus, the number of vulnerabilities also increases every year. This creates major challenges for companies, as they need to maintain an overview of vulnerabilities and affected systems in order to decide which vulnerabilities need to be addressed immediately which can wait or do not need to be addressed at all. Prioritization is important as vulnerability management resources are often limited.

Security advisories are a popular approach to provide this prioritization overview [10]. They contain information on vulnerabilities and affected products, and provide guidance on how to deal with the vulnerability. They are usually provided by vendors for their products and distributed by aggregators such as the NVD[1]. However, the security advisories are usually not available in a predefined format and the vendors themselves often decide how their security advisories are presented and distributed. Companies relying on solutions from various vendors often find themselves caught in a long chain of dependencies. The literature has shown that installing software updates is very time-consuming and often causes problems [8, 14], as each company uses various systems that use distinct products from different vendors, creating a complex and difficult-to-maintain environment. The challenge for a security analyst is to reliably find, filter and process information on software vulnerabilities in their organization. However, it is not known exactly how this process is carried out, what role security advisories play in it and what hurdles can be encountered. There are indications that the current state of security advisories is far from ideal. For instance, obtaining relevant information is often tedious due to the many irrelevant security advisories that are received [2, 10]. Sometimes essential information is also missing from an advisory, making it hard to fully understand the impact of the affected system to make an appropriate decision [7]. This implies that a standardized format with a uniform structure is needed.

[1]National Vulnerability Database, https://nvd.nist.gov/

The Common Security Advisory Format[2] (CSAF) is designed to enable automated processing of security advisories due to its standardized format. Nevertheless, a new standard will only be accepted by users if it solves existing problems and can easily be adopted. In order to draw a comparison as to whether CSAF can lead to an improvement, the current situation must be investigated. This raises the questions about how security advisories are currently processed, what problems can arise and whether CSAF provides a framework that addresses the existing challenges.

Contributions. We summarize the contributions of this work as follows:

- We shed light on how security advisories are currently received, processed and what influences the decision-making process. For example, advisories are usually received through one main channel, such as email. In addition, multiple mailing lists are often subscribed simultaneously due to the concern of missing relevant advisories. How a decision is finally made with the advisories depends heavily on the company policy, although the probability and potential impact of a vulnerability are consistently considered important.
- We highlight problems caused by the inconsistent format of security advisories. For example, inconsistent product identifiers defined by vendors complicate the process of identifying whether the advisory is relevant, as this information often has to be extracted manually. Thus, automatic pre-filtering is made more challenging, and non-relevant advisories are often received.
- We show that there is a great need and desire for a standardized format and automated processing. However, some users do not yet use automation, as adapting their systems to it is too resource-intensive.
- We show that CSAF can improve the processing of security advisories, as the machine-readable, standardized format enables automatic processing and saves resources. Nevertheless, solving some problems remains in the hands of the vendors, such as the completeness of the information contained in the security advisories. CSAF specifies a format, but not all entries in this format are mandatory.

Outline. The paper is structured as follows. First, in Section 2, we give insight into the background of security advisories and CSAF, outline related work and present the research questions. Next, in Section 3, we describe the qualitative preliminary study in which we conducted interviews. The design and results of the quantitative main survey is described in Section 4. The findings are then discussed in Section 5 and connected to related work and the real world. This work is concluded in Section 6, where we provide a brief summary and outlook for future research.

2 BACKGROUND

2.1 Security Advisories

Security advisories are documents that are usually issued by vendors consisting of valuable information about a vulnerability in one or more of their products. Sometimes, advisories are also issued by independent security researchers or cyber security organizations on a vulnerability they have discovered or researched. The main goal of these advisories is to inform users of the product, which are usually organizations, as well as other security researchers about these vulnerabilities in a comprehensive manner and also give recommendations on how to deal with them. This makes security advisories a very important tool for combating security threats which is why they are widely used today. As there was previously no recognized standard for them, security advisories were and still mostly are shared in various formats that range from simple text files to proprietary data types. Their content as well as its ways of dissemination are mostly dependent on the issuer. Still, they generally follow the same content structure which consists of a list of affected products, followed by details about the vulnerability and a recommendation for action. However, these sections differ a lot in detail and writing quality, depending on the author.

As far as the distribution of these documents is concerned, a variety of different methods in various combinations are applied. These include, but are not limited to, social media posts, blogs or mailing lists. Often users have to actively take care of obtaining security advisories. Some platforms or tools by different vendors also have made it their mission to simplify this process, by offering a service that collects, filters and distributes security advisories from various sources to interested parties.

2.2 CSAF

The Common Security Advisory Framework (CSAF) consists of a specification on how security advisory documents in the CSAF format are to be structured and which information they can contain as well as a set of tools that facilitate the creation, processing and distribution of them. The framework is created and maintained by OASIS[3], a well-known non-profit organization committed to developing open standards and aims to be the de facto globally known and used standard for generating security advisories. To achieve this, OASIS has the support of numerous large international institutions and corporations, such as Cisco, Oracle or Siemens, which were actively involved in the development process. CSAF documents are written in JSON[4] and are generally divided in three main sections. The document section contains various metadata of the document itself, such as the title, category and publisher details. The product tree section lists every single product that is referenced in the advisory as well as their relation to other products. It additionally can contain unique identifiers or identification helpers for them. The vulnerabilities section describes one or multiple vulnerabilities of the aforementioned products in great detail and also provides the reader with instructions on dealing with them. This is often accompanied by references to known vulnerability scores or descriptors, such as CWE[5], CVE[6] or CVSS[7]. Only a subset of the provided fields are mandatory for a CSAF document to be valid and a lot are optional. OASIS also thought about the distribution of CSAF documents and therefore put an infrastructure of issuers

[2]https://docs.oasis-open.org/csaf/csaf/v2.0/csaf-v2.0.html

[3]Organization for the Advancement of Structured Information Standards, https://www.oasis-open.org/

[4]JavaScript Object Notation, https://www.json.org/json-en.html

[5]Common Weakness Enumeration, https://cwe.mitre.org/

[6]Common Vulnerabilities and Exposures, https://www.cve.org/

[7]Common Vulnerability Scoring System, https://www.first.org/cvss/

and distributors of CSAF documents in place which allows users of the CSAF standard to quickly aggregate and filter advisories from multiple sources, depending on their needs.

2.3 Related Work

As CSAF is very new, there exists no current research on its possible effect on vendors and users. Security advisories are also not a particularly well-studied topic, although some work has examined them in different contexts. User studies about the update behavior of system administrators and developers proved to be very useful for our studies, since they investigated the handling of security advisories to a certain extent, even if these were not the focal point.

2.3.1 Automation. Fenz et al. [4] analyzed structures of existing security advisories with the aim of identifying standardized semantics to enable automated processing of advisories. The authors evaluated the advisories primarily in terms of semantic usefulness, information complexity and distribution. The results indicate that none of the existing formats met the authors' criteria. The paper from 2008 shows how long the desire for a standardized format has been under discussion. Building on this work, Fenz et al. [3] developed a framework that converts security advisories from different sources into a machine-readable format to enable automation. Ramnani et al. [12] focused on the format of security advisories and their automated processing. They used pattern recognition and natural language processing techniques to extract valuable information from large quantities of unstructured or semi-structured vulnerability information. For this purpose, they developed a prototype that they then evaluated on a test set. The authors emphasize the importance of automatic processing of security advisories.

2.3.2 Processing of Security Advisories and Updates. Li et al. [8] conducted a quantitative and qualitative study on how often and when system administrators perform updates. The authors identified five phases for update processing: 1) the system administrator is informed about the update and begins to gather information about it; 2) a decision is then made as to whether the update should be carried out or not; 3) the systems are prepared and the update is carried out, whereby 4) the system administrator communicates the update to the employees choosing a suitable time frame; 5) any problems that the update may have caused are resolved. The results indicate that the phases are quite complex and time-consuming for the system administrators. According to the authors, many of the steps could be optimized using automation. The same topic was addressed by Tiefenau et al. [14], who also interviewed and surveyed system administrators. They examined the update process and identified similar phases: information gathering, decision, update preparation, test run, post-installation. These first two steps – information gathering and decision-making – in the context of security related updates revolve primarily around security advisories and are of particular importance to our studies. Considering previous research on automation and the emphasis of CSAF on automated processing, we think that processing of gathered information (including security advisories) represents a crucial additional step between information gathering and decision-making. Using the insights of related work [8, 14], the processing of security advisories can be classified into three main steps: 1) information gathering, 2) information processing and 3) decision-making. We used these steps to structure the interview guide and the survey questionnaire. The steps are briefly explained in the following.

Information gathering. This step involves determining which channels are used and how the information reaches the user. Miranda et al. [10] analyzed security advisory platforms to determine how vulnerabilities are first published and distributed to other platforms. They found two types of platforms: information sources and aggregators. Vulnerability information spreads through a network of connected platforms. To investigate if developers keep software libraries up to date, Kula et al. [7] examined various GitHub projects and found many outdated dependencies which contained security vulnerabilities. When inquired about this issue, developers often replied that they were not aware of the outdated dependencies. Most of them stated that they simply do not have the time to search for corresponding security advisories and generally lack resources to update their dependencies. Farhang et al. [2] analyzed various bulletins which vendors often publish for their own platforms, such as Android, to summarize security-relevant events. The authors show that the vendors often do not use the standardized CVEs, but their own identifiers in the bulletins which makes it unclear whether all critical vulnerabilities are named and whether the information is relevant. These studies show that a main channel to receive advisories would be desirable, as this would make the process less complex. This demonstrates that it is highly complex for security analysts to receive all relevant information. Often, all sorts of security advisories are received, some of which are not relevant [10]. Another problem is that the information is often distributed across different sources and users have to tediously gather and filter it [7, 10]. Trustworthiness of the sources also plays a role [6], which CSAF ingrained in their design, as CSAF trusted providers have to sign and hash their advisories [11].

Information processing. Once information has been collected, it is processed either manually or with the help of tools [10]. If tools are used, it is particularly important that the information is available in a complete and machine-readable format, as otherwise it cannot be processed [3, 4]. The absence of a standardized format means that the information often appears redundantly in the security advisories or is widely distributed across different sections, such that filtering takes a lot of time. Users would therefore like to have better options for filtering the information with the help of a tool [10].

Decision making. The final step is to decide how to deal with the vulnerability. It is particularly important that the vulnerability is described in as much detail as possible in the security advisories to fully understand its impact on the affected systems [7]. Users would like to have more opportunities to contact other affected users or the vendors directly in order to simplify the decision-making process [7]. The exact procedure to be followed afterwards depends on company policies and whether a patch would threaten the current availability of systems or resources [8, 14].

2.4 Research Questions

Security advisories play an important role in vulnerability management. Therefore, it is important to know how exactly users

work with security advisories and which steps they go through when processing them. For example, from which sources they receive advisories, what information is relevant and which actions are taken.

The literature as described in Section 2.3 indicates that the processing of security advisories is very time-consuming and complex. Possible problems may arise at various points in the processing workflow and further tie up valuable resources. We thus investigate which problems may occur based on the current state of security advisories.

An important feature of CSAF is the ability to process security advisories automatically. However, the precondition for this is that users are willing to use automation and to adapt their own systems and setups if necessary. We shed light on the question of how automation is currently used for security advisories and what the users' intentions are, for example, whether they are already using automation or are planning to introduce it soon.

One finding from the literature was that problems often arise due to the inconsistent structure of advisories. CSAF attempts to tackle this problem by creating a standardized and machine-readable format that can be processed automatically. This raises the question to what extent CSAF is in current use, and which other problems CSAF plans to solve with its current approach.

To summarize, we consider the following research questions:

RQ1: How are security advisories currently processed?
RQ2: Which issues do security advisories cause in their current state?
RQ3: What is the current state of automation for processing security advisories?
RQ4: Which security advisory issues can CSAF address and which not?

The research questions are investigated using a qualitative preliminary study and a quantitative main survey. RQ1-RQ3 are examined directly by the results in Section 4.3. RQ4 is evaluated by the interpretation of the results in Section 5.

3 PRELIMINARY STUDY - INTERVIEWS

3.1 Study Design

Interview Guide and Testing. Three qualitative interviews were conducted to identify findings that would be further investigated in the quantitative survey. The interview guide was developed from similar topics previously found in the literature review. The aim of the preliminary study was to discuss the problems that were identified during literature review in more detail. This was done to get a rough estimation of how widespread and severe they are perceived, and to add other aspects that were not previously considered, such as the fear of missing relevant security advisories.

The interview guide can be found in Appendix A. First, the interview participants were asked to introduce themselves and give a brief overview of their workplace and work experience, for example how often they come into contact with security advisories. Next, they were queried about collecting information, for example from which sources they obtain advisories, whether they receive too many irrelevant advisories or if there are any other problems. This was followed by questions about the processing of security advisories, such as whether they are processed manually or automatically. We requested the participants to name the advantages and disadvantages of automation or manual work that is currently used in their day-to-day work. We then went on to discuss the decision-making process once the information from the security advisories had been processed. Among other things, they were asked which factors can influence a decision. Finally, the interviewee's demographic data was recorded and room was provided for further comments. To ensure that the questions were comprehensible and to estimate the duration of the interviews, a test interview with a colleague from our lab was conducted.

Ethics. The data protection office of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) approved the study. Before the recording started, the participant was provided with an informed consent form that informed them about their rights of information and deletion, and the purpose of the study. There was no compensation and the participants could interrupt the interview at any time and decline to answer. During and after the interview, it was possible to let us know if certain information should not appear in the final transcript. The answers were pseudonymized and the recordings stored on a server of our university. All interviews were conducted remotely via Zoom.

Recruitment and Data Analysis. We focused on recruiting participants who work intensively with security advisories on a day-to-day basis and are familiar with CSAF to inquire about CSAF-specific topics, such as automation. To advertise the interviews and the subsequent survey, we gave a short talk about the project at a conference for security professionals in Germany[8] . This way, we recruited two interview participants and a third participant was recruited through personal contacts from our lab. The interviews took 22 minutes on average. They were recorded and then transcribed manually by one researcher. All personal references were removed in the process. After that, the transcript was analyzed by one researcher in order to fill previously identified themes of security advisory processing from literature with content. For example, the exact procedure and possible problems were collected for each processing step. Due to the small sample size and the fact that the broad categories (information gathering, information processing, decision-making) had already been determined after the literature review, it was decided not to use coding. While reviewing the transcripts, new themes that had not previously appeared in the literature were added to the list of themes, such as the willingness for automation from a user perspective. After the interviews were analyzed, the identified themes were discussed by two researchers to draw conclusions.

Participants. All three participants come from Germany, are male, have a university degree and were on average around 55 years old at the time of the interview. One of the participants works as head of IT, one is a system administrator. The third participant works as a service manager at a company that publishes security advisories. All participants have many years of professional experience with security advisories and vulnerability management.

[8] https://www.dfn.de/event/30-dfn-konferenz-sicherheit/

3.2 Results

Processing of Security Advisories. To receive security advisories, the interviewees use the DFN[9] service which informs subscribers by email about vulnerabilities that affect them. Subscribers are also informed when something has changed in the vulnerability which the interviewees emphasized as particularly positive. However, one participant noted that he was unable to change the settings of his subscription and consequently also receives some security advisories that do not apply to his systems. The other participants also said that they occasionally receive security advisories that do not match their current system or setup.

One problem stated by the interviewees is a mismatch in the product identifiers which normally mark products with a unique code. According to the participants, there are often several different product identifiers for the same product, as companies assign their own identifiers internally. It is thus possible that various product identifiers are specified, but they refer to the same product. It is then necessary to manually check which products are actually affected. This is described by the participants as very time-consuming, as the list of affected systems is often very long and in an inconsistent format, usually consisting only of a sequence of letters and numbers. This makes it particularly difficult to determine whether one's own system is affected by the vulnerability or not. P3 summarized this problem as: *"If we write in S7-1511 and the customer wrote Sematic S7-1500 Family or something like that, [...] even if he has an asset management system and we have machine-readable advisories, the matching is — you could try it with AI or something — [...] not that simple. In other words, the issue of unique product identifiers is still unresolved throughout the industry."* The interviewees consider automation to be difficult without a standardized format for product identifiers, as they fear that they will miss important advisories.

As described in Section 2.1, security advisories are usually provided by the manufacturers as PDFs or other text documents. According to the participants, another problem is the varying length and quality of the security advisories. For example, some advisories are too short, making it necessary to find additional sources. One participant noted that a certain manufacturer only offers links in its security advisories. Thus, the participants prefer the security advisories to be more detailed. In P1 opinion, *"there can't really be too much [information], because if a document is well-structured, it's not a problem"*.

The participants see a standardized format for security advisories as very positive. According to the participants, the way the vulnerability is subsequently handled depends on its severity and internal company policies. For instance, whether it requires immediate fixing or whether a patch cycle should be waited for.

CSAF. P1 and P2 are familiar with CSAF, but do not currently use it. P3 is an employee of a company that has cooperated with OASIS to publish security advisories in the CSAF format. Currently, the security advisories are already processed semi-automatically according to the participants by pre-filtering them depending on the affected system. Although the participants stated some problems, such as mismatches in product identifiers and the cumbersome manual research of which systems are exactly affected, the interviewees do not see a significant need for further automation. P3, who works as a service manager and therefore deals with customers who process security advisories on a daily basis, also emphasized that many users are currently not ready for automation. This is because it requires a functioning asset management system that is always up to date, and lists products and setup details, which many companies do not have yet. The interviewee stated: *"If you consume machine-readable advisories, you actually need an asset management system so that you can manage the matching. And that is usually lacking on the customer side."*

Overall, the interviews revealed that many of the problems with security advisories arise due to an inconsistent format. Nevertheless, automation is still viewed skeptically at present, as companies would have to adapt their current setup to automation.

4 MAIN STUDY — SURVEY

4.1 Study Design

Questionnaire and Testing. Similar to the interviews, the main part of the survey was structured around the three parts, 1) information gathering, 2) information processing and 3) decision-making, which we identified from the literature review. The content and wording of the questions was derived from the findings of the preliminary study. The full questionnaire can be found in Appendix B. The language of the survey was English, as we aimed to recruit as many participants as possible. First, the participants were asked whether they regularly work with security advisories. As we are focusing specifically on this target group, the survey was ended at this point for all participants who do not handle security advisories. The participants were then queried about their occupation and general IT security expertise, such as in which economic sector they are working, their profession and the size of their company. Next, we asked questions about how frequently security advisories are dealt with and familiarity with CSAF. This was followed by the first major part about information gathering, were we asked which types of channels the participants receive their advisories, such as email or messengers. We then provided statements concerning the reception and sources of advisories, such as "I receive security advisories that do not affect me", which participants could agree or disagree with using a 5-point Likert scale. These statements were derived from the problems identified during the preliminary study.

Next came questions about the processing of security advisories, such as whether an asset management system or automation is used. This was also followed by statements, originating from the preliminary study, the participants could agree or disagree with, this time concerning the processing of security advisories, such as "Automation is essential to handle security advisories efficiently". An attention test was also included here. After this, the participants were asked questions concerning decision-making, e.g., if they are involved in this process and how important are some factors, such as company policies and available resources.

Finally, the participants were given an opportunity to comment in a free text field if something was not covered by the previous questions. The survey was concluded with demographic questions.

We conducted five incremental test runs, each with four to five testers. After each test round, the collected feedback was integrated.

[9]Deutsches Forschungsnetz (German research network), https://www.dfn-cert.de/

The testers were mainly IT security researchers of our lab and institutes who research or work in the IT security context. For the last round, we recruited people from the industry who work with security advisories on a daily basis.

Ethics. The study was approved by the data protection office of Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). At the beginning of the survey, the participants were informed about their data protection rights and the purpose of the study in an informed consent form. Participants were only able to take part in the survey if they gave their consent. There was no compensation, participation was voluntary. In addition, participants could interrupt or cancel the survey at any time without any disadvantages. The demographic questions were optional. An "Prefer not to say" option was offered and preselected for all demographic questions. All responses were anonymized. The survey was carried out on a self-hosted LimeSurvey[10] instance at our university.

Data Analysis. For the quantitative analysis, the data was first extracted from LimeSurvey as JSON files and then analyzed using Python scripts. For correlation analysis, χ^2 was used as the variables are categorical, and the effect sizes were determined by Cramér's V [1]. Considering the qualitative analysis, we followed the exclusion criterion of McDonald et al. [9] as "coding requires little interpretation" [9, p. 72:3]. Since there were only 20 comments in total, and these were brief and clearly understandable, e.g., "The quality of writing in a security advisory is important" or "I consider qualified automatic actions on reported threats to be completely impossible", we decided not to code them. One researcher summarized the free text answers. They were then discussed by two researchers to decide on the final findings and used to illustrate the results in the following.

Recruitment. In order to advertise the online survey for international participants, we first posted the survey on fitting Reddits[11] and used personal contacts who spread the survey on LinkedIn. We also advertised the survey using IT security mailing lists such as Fulldisclosure[12] and Blueteamsec newsletter[13]. With the assistance of DFN-CERT, the link to the survey was distributed to the DFN community.

4.2 Participants

The survey ran from September 2023 to February 2024 and 230 people completed it. Some participants were excluded as they did not agree to the privacy policy or did not pass the attention test, resulting in a total of 197 valid responses.

An overview over the participants' demographics is presented in Table 1. The participants were on average 44 years old (median 45 years). About 80% identified as male, 3% as female, 1% as diverse and 16% preferred not to disclose this information. About half of the participants come from Germany, 18% come from another European country and 8% from the USA. Regarding education, 48% of participants have a Master's degree and around 24% have a Bachelor's degree. Only about 12% of the participants have no academic or similar degree.

	N	%
Total	197	100,0
Male	158	80.2
Female	6	3.0
Diverse	2	1.0
N/A	31	15.7
18-29 years	11	5.6
30-39 years	41	20.8
40-49 years	44	22.3
50-60 years	37	18.8
Above 60 years	10	5.1
N/A	54	27.0
Germany	103	52.3
Other European country	35	17.8
USA	16	8.1
Other country	8	4.1
N/A	31	15.7
No academic education	24	12.2
Bachelor's degree	47	23.9
Master's degree	94	47.7
Ph.D.	9	4.6
Other	4	2.0
N/A	20	10.2
1-5 years work experience	27	13.7
6-10 years work experience	26	13.2
11-20 years work experience	60	30.5
More than 20 years work experience	84	42.6
None or basic IT security knowledge	1	0.5
Intermediate IT security knowledge	25	12.7
Advanced or expert IT security knowledge	171	86.8

Table 1: Overview of the participants' demographics, work experience and self-assessed IT security knowledge.

Most of the participants work in the IT field and have a job title that is directly related to IT security, such as "security analyst", "cyber security specialist", "pentester" and others. Other participants described their job title as "head of IT", "IT lead" or "coordinator", which means that they work as a leader or supervisor of a team and are presumably responsible for making decisions. The remaining participants were mostly system administrators or developers.

On average, participants had 19 years of professional experience (median 20) and most considered their expertise in IT security to be advanced (42%) or expert (45%). Most participants work in the education sector (25%) or information and communication services (23%). The size of the companies was fairly evenly distributed. About 10% of the participants work in a company with less than 50 employees, 17% in one with 50-249, 16% in one with 250-999, 36% in one with 1,000-9,999, and 23% in companies with over 10,000 employees. Over 90% of participants engage with security advisories on at least weekly basis, 9.1% receive them monthly. Around 85% of participants forward security advisories at least sometimes to potentially

[10] https://www.limesurvey.org
[11] For example https://www.reddit.com/r/AskNetsec
[12] https://seclists.org/fulldisclosure/
[13] https://bluepurple.binaryfirely.com

affected end customers and industry partners, and about 93% of participants are at least partially involved in the decision-making process.

4.3 Results

Information Gathering (RQ1, RQ2). Only about 12% of the participants use a single channel, most use about 2-3 channels. Receiving security advisories per mail, web platforms and via colleagues are most popular, while messengers are less common. Other channels that were stated as free text answers were RSS feeds and social media, such as Mastodon or X, but also blogs, podcasts or company internal software. On average, participants are subscribed to 12 different mailing lists or similar services. We consider using different mailing lists as one channel.

We presented a list of statements to the participants that they could agree or disagree with using a 5-point Likert scale. We coded the Likert scale from one (strongly disagree) to five (strongly agree) to calculate the mean and standard deviation, as recommended by Harpe [5]. The results are depicted in Figure 1. Many participants frequently receive security advisories that do not affect them. In addition, the participants do not have the impression that they are missing out on important advisories. This is probably due to the fact that most participants obtain their information from several sources and have subscribed to multiple services. As a result, they receive more advisories that are not necessarily relevant to them, but do not miss the particularly important ones. The advisories are sent to the participants shortly after they are published. The participants thus do not encounter problems with delayed information. The question whether the participants trust their sources that provide them with security advisories received rather mixed opinions. Although the majority agrees, there are also participants who do not trust all their sources. Many participants see the need for a central register for security advisories, but 25% of the participants were undecided. P59 stated: *"With multiple sources of security advisories available there is definitely a need for a standard format that is machine readable [...]"*.

Information Processing (RQ1, RQ2). Around 60% of participants use an asset management system to maintain their companies' or organizations' IT inventory. At least 20% would like to introduce one.

Around 26% describe themselves as familiar with CSAF, while the majority are only slightly familiar (40%) or not familiar at all (34%). This result is also reflected in the use of automation, as only about 33% of participants already use automation and about 25% have the desire to introduce it. However, around 42% of participants do not want to use automation. We dive deeper into usage of automation in Section 4.3. Considering the personal opinion about automation, most participants see automation as essential to handle security advisories in an efficient way (see Figure 3). They also clearly see the need for security advisories to have the same structure. The participants also agree that organizations need established vulnerability management procedures in order to process security advisories. In contrast, there is mixed agreement on how easy it is to match the affected systems listed in a security advisory to one's own IT setup. Around 24% of participants think it is easy to identify the required information, while 43% find it complicated and around 27% were undecided. This confirms the finding from the qualitative study, in which the participants reported that the inconsistent format often made it very difficult to categorize the flood of affected systems and to identify whether one's own system is affected or not. When answers to this statement are compared with the size of the companies (see Figure 2), it is noticeable that participants who work in companies with 50-250 employees find matching easier than participants of larger and smaller companies. This could be because smaller companies often have a less complex IT infrastructure compared to larger ones, but have more resources to do matching than the smallest companies of up to 49 employees.

Decision Making (RQ1, RQ2). After the important information has been extracted from the security advisory, a decision must be made on how to proceed with the vulnerability. Some factors play a more prominent role than others (see Figure 4). According to the participants, the most important factor is the likelihood that the vulnerability will be exploited and the potential impact of exploitation. The internal company policies and the available resources have a slightly smaller impact, although these are nevertheless considered to be important or very important by at least 50% of the participants. If no decision can be made on the basis of the information available in security advisories, additional sources must be consulted. About 70% of the participants use at least some additional sources to reach a decision. We asked these participants to describe their additional sources in more detail in a free text field. These sources are often public security scores that provide an approximate security assessment, such as CVSS, or databases from manufacturers and vendors that provide further information such as patch notes. Software used for vulnerability scanning or management is also utilized. Sometimes additional expertise is obtained from security consultants, peers or colleagues, or publicly available information is used, such as news sites, blogs or social media. In rare cases, internal company records of past incidents, decisions or best practices are also used.

Automation (RQ3). About 33% of participants already use automation and about 25% have the desire to introduce it. However, around 42% of participants do not want to use automation. If the use of automation is correlated to familiarity with CSAF (see Figure 5), it seems that people who use automation are significantly more likely to know about CSAF ($\chi^2 = 18.19$, $p < 0.05$). Cramér's $V = 0.215$, which can be interpreted as small to moderate effect size.

Another finding is that automation seems to be significantly related to the size of the company, see Figure 6 ($\chi^2 = 18.34$, $p < 0.05$). Cramér's $V = 0.216$, which indicates a small to moderate effect size. Larger companies with more than 10,000 employees tend to use automation more than smaller companies. This is not surprising, as larger companies often have more resources.

The use of asset management systems also appears to be significantly correlated to the use of automation ($\chi^2 = 18.56$, $p < 0.05$), see Figure 7. The effect size ranges from small to moderate (Cramér's $V = 0.217$). Most likely, the use of an asset management system simplifies the introduction of automation.

Additional remarks. At the end of the questionnaire, we gave the participants the opportunity to express further thoughts on the topic in the form of a free text field. Some participants again

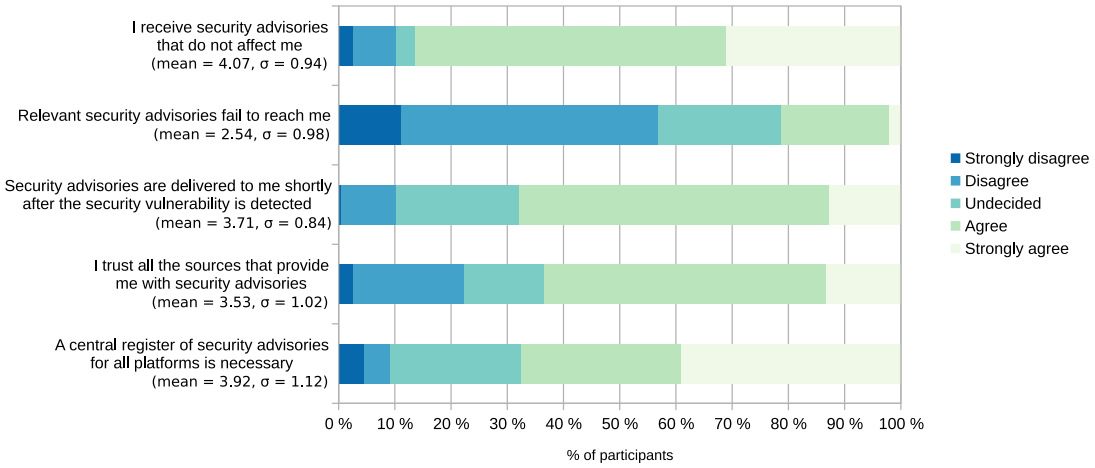

Figure 1: Statements and participants' agreement for information gathering ($N = 197$).

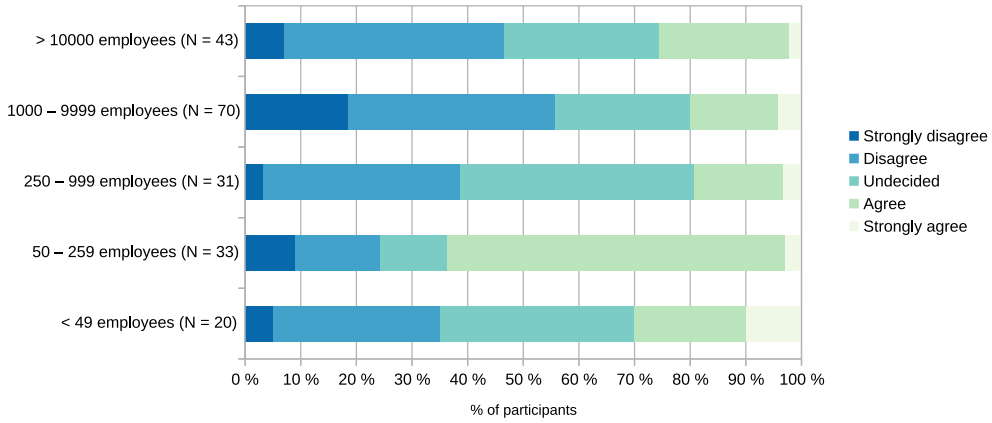

Figure 2: Agreement for "It is easy to match the affected systems of a security advisory to our IT setup" in relation to company size ($N = 197$).

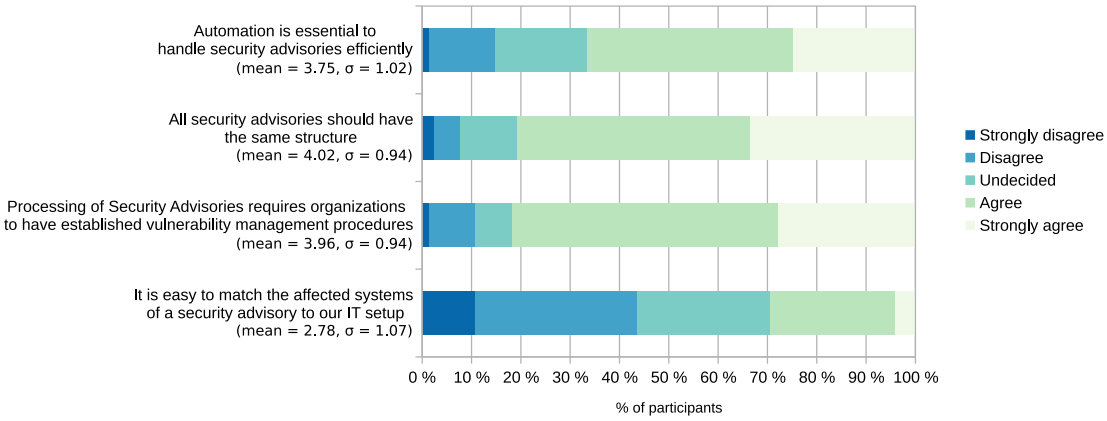

Figure 3: Statements and participants' agreement for information processing ($N = 197$).

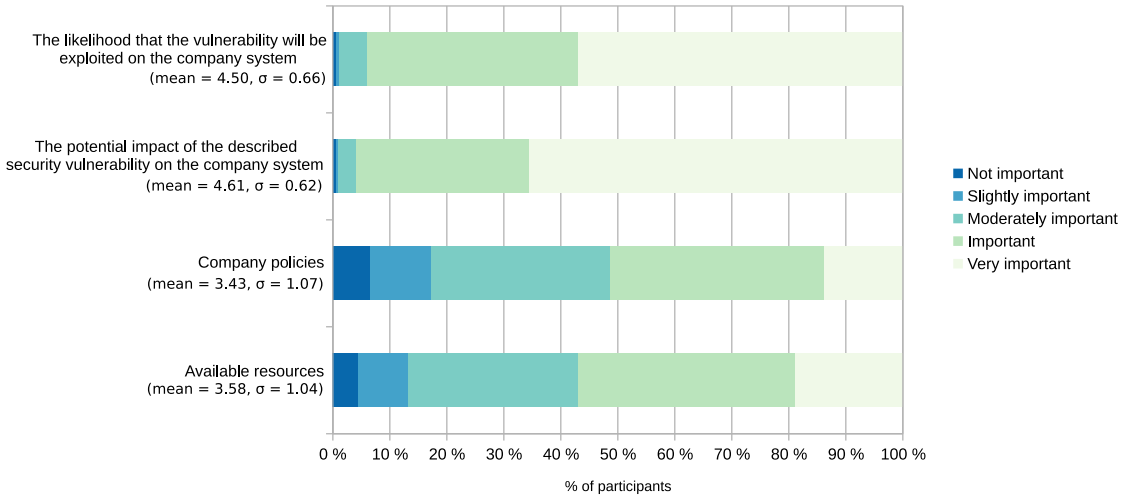

Figure 4: Factors that may influence decision-making regarding the security issue described in an advisory, and their importance to the participants ($N = 197$).

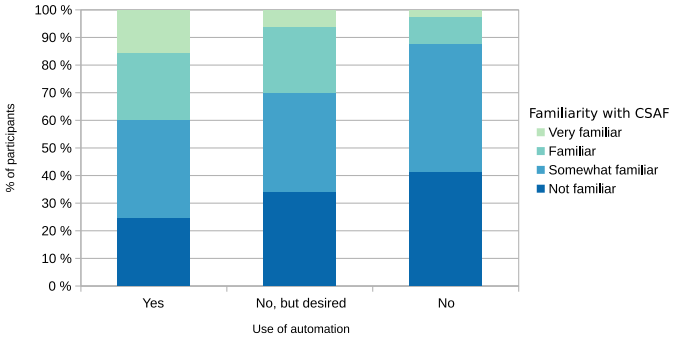

Figure 5: Use of automation in relation to CSAF familiarity ($N = 197$); $\chi^2(6) = 18.19$, $p < 0.05$, **Cramér's** $V = 0.215$.

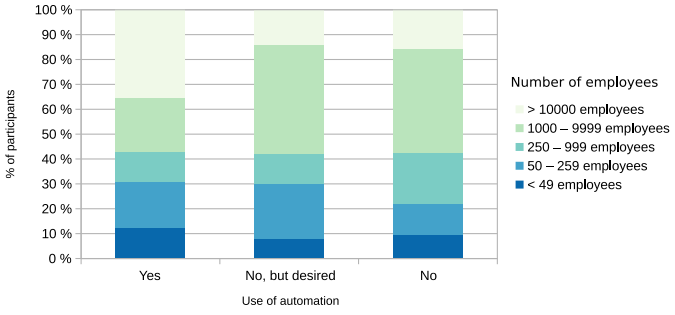

Figure 6: Use of automation in relation to company size ($N = 197$); $\chi^2(6) = 18.34$, $p < 0.05$, **Cramér's** $V = 0.216$.

emphasized the poor quality of security advisories, which manifested itself in poor writing quality, bad wording or wrong product and version numbers. *"Security advisories often contain wrong data,"* wrote P51, and P14 stated: *"People must learn how to write."* Other

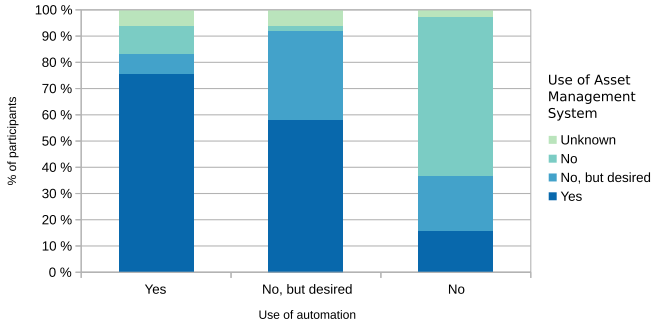

Figure 7: Use of automation in relation to usage of asset management systems ($N = 197$); $\chi^2(6) = 18.56$, $p < 0.05$, **Cramér's** $V = 0.217$.

participants used this section to criticize existing security scores such as CVSS. Some participants also emphasized once again that they are skeptical about automation, as P137 described: *"I consider qualified automatic actions on reported threats to be completely impossible"*. P97 explained this further: *"A fully automated process to install security updates is not feasible for all used products. Especially propriatary [sic] software [...] have a worse testing in place before releasing a software update and might break a running system. In recent years this problem has worsened a lot as the vendors internal policy seems to reduce the spending on testing and internal reviews of their products."* P131 stated: *"Vulnerability management is such a complex challenge that it generally lacks ressources [sic]."*

5 DISCUSSION

5.1 Current State of Advisories (RQ1, RQ2)

The processing of security advisories can be divided into different phases, as it was identified using related work [8, 14]. First, advisories must be received from sources. They must then be filtered

according to relevance. Then a decision has to be made on how to proceed with the vulnerability. However, little was known about how exactly these phases are structured and what role security advisories play in them which was investigated through this study. The results showed that security advisories are processed on an almost daily basis which requires a high level of resources. In most cases, several sources are used simultaneously, as one source is often not sufficient to obtain all the information required. The participants therefore see a need for a central source that distributes the security advisories. Due to the different sources, many security advisories are also received that are not applicable for the users at all. In this case, users must first go through the time-consuming step of manually analyzing the security advisory, and determining whether their own product is affected which is very cumbersome due to the inconsistent format. In addition, the manufacturers often do not use the same unique product identifiers in the security advisory, but internal identifiers, which Farhang et al. [2] also noted, though they did not focus on the user side. This makes this step unnecessarily resource-intensive and complicated.

We found evidence that the participants see the need for a standardized format with a uniform structure, as it would solve this problem. The decision-making process then depends on the company's objectives, with the likelihood of exploitation and the impact of the vulnerability being particularly important according to the participants. Additional sources are often used in the decision-making process which in turn could be linked or integrated into security advisories in order to save additional resources.

5.2 Automation (RQ3)

Automation could save additional resources and simplify time-consuming processes, which is why automation is emphasized as particularly important [3, 4, 12]. Nevertheless, some participants are not currently planning to switch to automation. This may be mainly due to the fact that some prerequisites must first be met for automation to be compatible. For example, an integrated asset management system must be in place, maintained, and managed to provide an up-to-date list of product numbers and current version numbers. Introducing such a tool into an existing system can involve a great deal of effort, for which current company resources may not be sufficient. Some participants are skeptical about automation, as they believe that automation cannot solve all problems, such as the problem of inconsistent product identifiers.

5.3 Need for CSAF (RQ4)

CSAF, as a standardized format, would solve some of the problems mentioned above. It is likely that users will continue to receive irrelevant advisories, as the problem of inconsistent product identifiers is in the hands of the manufacturers. However, with CSAF it would be easier to manually find out whether users' systems are affected, or this step could be even fully automated. More information could also be added to security advisories, which the participants of this study would welcome, as they believe that there cannot be too much information in a security advisory. Nevertheless, many information entries in CSAF are optional, and it is up to the manufacturers and vendors to decide how much information they provide and how well written it is. The main criticism from participants is that some manufacturers pay too little attention to their advisories, and that the general quality of advisories is rather poor.

CSAF is the right way to achieve automated processing of security advisories and to simplify complex processes in order to save valuable resources. Nevertheless, there are some problems that even a standardized format cannot solve. If information fields are optional, it is up to the manufacturers and vendors to decide how carefully they enter this information, and what its quality is. In addition, an automated format only makes sense if users are prepared to accept automation and adapt their own systems to it which was not always the case in our study.

5.4 Limitations

Internal Validity. The survey aimed to investigate how people process security advisories, so we recruited individuals who regularly encounter security advisories using a filter question. We avoided mentioning CSAF in the beginning of the survey to include both familiar and unfamiliar participants. We conducted multiple rounds of testing with various participants, including four industry professionals, to ensure that the questionnaire was understandable and purposeful to our research questions. Participants were mainly asked for facts, not opinions, to minimize potential question order effects. Although we couldn't confirm if participants took the survey more than once, we believe that this is unlikely, as we offered no compensation.

External Validity. This study might be limited by the small sample size of the qualitative preliminary study, as only three people were interviewed. However, as the interviews served to gain an initial insight and collect problems that would then be examined in more detail in the survey, we did not expect that the interviews to be representative of the community who deals with security advisories regularly. Another limitation is that slightly more than 50% of participants of the main survey are from Germany, though efforts were made to recruit internationally. With participants representing a geographically more diverse group, different conclusions might have been reached. We took the utmost care to find out important topics and issues concerning security advisories through literature review and three interviews. Nevertheless, this survey might have left out some aspects of working with security advisories if they did not come up in the preparatory work. Due to the survey's quantitative nature, only the frequency of identified aspects can be described, and empirical evidence of the reasons behind them cannot be provided.

6 CONCLUSION

In this work, we investigated how security advisories are currently used, where problems arise and whether CSAF can improve this process.

First, we conducted a qualitative preliminary study in which we interviewed three participants about their handling of security advisories. A particular focus here was on the establishment of automation and the participants' opinions on this, as CSAF emphasizes automation. We then conducted a quantitative online survey with 197 participants. The results show that users encounter different problems when processing security advisories, and that these

problems occur at different phases. For example, participants often receive many irrelevant security advisories and then have to filter them manually. This is particularly resource-intensive, as the version numbers concerned are often in a non-uniform format and confusing. CSAF could solve this problem with a standardized format. Many participants see the advantages of CSAF and automation, and recognize the need for improvements. Nevertheless, there are hurdles to overcome as some users are not yet fully convinced by automation or lack the prerequisites for its use.

As over half of the participants came from Germany, future work could extend the research to other countries. Further studies could also explore manufacturers opinions on CSAF, including their process for detecting potential security flaws and creating security advisories. Another follow-up study could evaluate the outcomes of this research. In this case, the results could be discussed with CSAF users to identify further trends. Alternatively, the findings could be discussed with the developers of CSAF in order to potentially contribute to improvements.

This work showed that CSAF is not yet widely used. If it becomes more widespread in the future, a study could explore whether CSAF was actually able to solve the problems identified in this study and improve the processing of security advisories.

Acknowledgements. We thank Martin Waleczek, Tobias Dussa and Thomas Schreck for their valuable support in recruiting participants. We also thank Freya Gassmann, who conducted the correlational statistical analysis. We thank the testers and participants of the survey. We thank the anonymous reviewers who improved the paper with their feedback. The authors were supported by the German Federal Ministry of Education and Research under grant 16KIS1271K.

Author Contribution Statements. **Julia Wunder:** Conceptualization, Methodology, Formal analysis, Visualization, Supervision, Project administration, Writing - Original Draft (Chapters 1, 2.3, 2.4, 3, 4, 5, 6), Writing - Review & Editing, **Janik Aurich:** Conceptualization, Methodology, Formal analysis, Investigation, Data Curation, Writing - Original Draft (Chapters 2.1, 2.2), Writing - Review & Editing **Zinaida Benenson:** Visualization, Supervision, Project administration, Writing - Review & Editing, Funding acquisition

REFERENCES

[1] Harald Cramér. 1999. *Mathematical Methods of Statistics.* Princeton University Press.
[2] Sadegh Farhang, Mehmet Bahadir Kirdan, Aron Laszka, and Jens Grossklags. 2020. An empirical study of Android security bulletins in different vendors. In *Proceedings of The Web Conference 2020.* 3063–3069.
[3] Stefan Fenz, Andreas Ekelhart, and Edgar Weippl. 2008. Fortification of IT security by automatic security advisory processing. In *22nd International Conference on Advanced Information Networking and Applications (aina 2008).* IEEE, 575–582.
[4] Stefan Fenz, Andreas Ekelhart, and Edgar Weippl. 2008. Semantic potential of existing security advisory standards. In *Proceedings of the FIRST 2008 Conference-Forum of Incident Response and Security Teams.*
[5] Spencer E Harpe. 2015. How to analyze Likert and other rating scale data. *Currents in pharmacy teaching and learning* 7, 6 (2015), 836–850.
[6] 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).* IEEE, 105–119.
[7] Raula Gaikovina Kula, Daniel M German, Ali Ouni, Takashi Ishio, and Katsuro Inoue. 2018. Do developers update their library dependencies? An empirical study on the impact of security advisories on library migration. *Empirical Software Engineering* 23 (2018), 384–417.
[8] Frank Li, Lisa Rogers, Arunesh Mathur, Nathan Malkin, and Marshini Chetty. 2019. Keepers of the machines: Examining how system administrators manage software updates for multiple machines. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019).* 273–288.
[9] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. 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 (2019), 1–23.
[10] Lucas Miranda, Daniel Vieira, Leandro Pfleger de Aguiar, Daniel Sadoc Menasché, Miguel Angelo Bicudo, Mateus Schulz Nogueira, Matheus Martins, Leonardo Ventura, Lucas Senos, and Enrico Lovat. 2021. On the flow of software security advisories. *IEEE Transactions on Network and Service Management* 18, 2 (2021), 1305–1320.
[11] OASIS. 2022. Common Security Advisory Framework Version 2.0, 7.2.3 Role: CSAF trusted provider. https://docs.oasis-open.org/csaf/csaf/v2.0/os/csaf-v2.0-os.html#723-role-csaf-trusted-provider Accessed in July 2024.
[12] Roshni R Ramnani, Karthik Shivaram, and Shubhashis Sengupta. 2017. Semi-automated information extraction from unstructured threat advisories. In *Proceedings of the 10th Innovations in Software Engineering Conference.* 181–187.
[13] Statista. 2023. Number of Internet of Things (IoT) connected devices worldwide from 2019 to 2023, with forecasts from 2022 to 2030. https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/ Accessed in March 2024.
[14] Christian Tiefenau, Maximilian Häring, Katharina Krombholz, and Emanuel Von Zezschwitz. 2020. Security, availability, and multiple information sources: Exploring update behavior of system administrators. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020).* 239–258.

A INTERVIEW GUIDE OF THE PRELIMINARY STUDY

A.1 Introduction

- Greeting, short introduction of people performing the interview + taking notes
- Short outline of thesis topic and why we're conducting interviews in the first place
- Notice on video / audio recording + approval by the interviewee
- Thank interviewee for their time and help

A.2 General demographic information

- Age, gender, country of residence, education
- Employment
  - job title, employment form
  - general routine / tasks

A.3 Security advisories general

- How often security advisories come up during work
- General view on security advisories in their current state *(positive, neutral, negative)*

Give an outlook of the following three categories and get across that they are to be analyzed one by one independent of each other.

A.4 Information gathering

Explain that this is only about collecting information.

- How are security advisories received
  - actively searched for
  - notifications
- Is there anything that makes gathering information particularly hard?
  - volume
  - sources
  - availability

- Is there anything you can think of that would improve the information gathering process?
  - time
  - security

A.5 Information processing

Assume all the necessary information is now collected. Clarify that this is only about the presentation of information and not about comprehension.

- In general, how is information processed
  - manual reading / skimming
  - automated work
- Are there any obstacles that make processing advisories difficult for you?
  - language
  - redundancy
  - complexity
- Is there anything you would wish for that would optimize information processing?

A.6 Decision making

So now the information is processed, what happens next.

- What factors are most influential in forming a decision on how to proceed?
  - perceived severity
  - available resources
- Under what conditions is it impossible to utilize security advisories to make an informed decision?
  - work environment / policies
  - security advisory content
  - general comprehension of issue
- Is there anything you would change or add to security advisories that would simplify the decision-making process?
  - risk score
  - contact information

A.7 Optional

Anything to add?

A.8 Conclusion

- Thank for the interview, ensure that their participation is of great value to the study

B QUESTIONNAIRE OF THE MAIN SURVEY

This is the questionnaire of the main survey. In the beginning, the questionnaire informs the participant about the purpose of the study, which is followed by the participation consent.

B.1 Preselection Question

Security advisories are used to communicate a security vulnerability in a system. They are usually issued by the manufacturer of the affected system and contain detailed information about the vulnerability as well as how to further proceed to mitigate the risk of being attacked.

- Do you regularly encounter security advisories as part of your work?
  - Yes / No

(If "No" is selected, the survey ends with a short explanation that we seek participants who regularly encounter security advisories as part of their work.)

B.2 Occupation and general IT security

- In which economic sector are you currently working?
  - Accommodation and food services / Administrative and support service activities / Agriculture, forestry and fishing / Construction / Distributive trade sector / Education / Electricity, gas, steam and air conditioning supply / Information and communication services / Manufacturing / Mining and quarrying / Professional, scientific and technical activities / Real estate activities / Repair of computers and personal and household goods / Services (except transport and storage) / Transportation and storage services / Water supply, sewage, waste management and remediation activities / Other: (free text)
- What is your profession? (free text)
- How large is the company you work for?
  - < 49 employees / 50 - 249 employees / 250 - 999 employees / 1000 - 9999 employees / > 10000 employees
- How many years of work experience do you have?
- Rate your level of knowledge in IT security.
  - None / Basic / Intermediate / Advanced / Expert
- How frequently do you engage with security advisories?
  - Daily / Weekly / Monthly / Yearly
- Common Security Advisory Framework (CSAF) is a standardized format for disclosing security vulnerabilities. It is designed to enable efficient creation and handling of security advisories with a focus on automation. How familiar are you with CSAF?
  - Not familiar / Somewhat familiar / Familiar / Very familiar

B.3 Sources for security advisories

- Please indicate the types of channels through which you receive security advisories. (multiple-choice)
  - Email / Messenger (SMS, Whatsapp, …) / Colleagues / Web platform / Search Engine / Other: (free text)
- How many different mailing lists or other services that provide security advisories are you subscribed to?
- Below is a list of statements about receiving security advisories. Please answer to what extent you agree with the following statements:
  (Strongly disagree, Disagree, Undecided, Agree, Strongly agree)
  - I receive security advisories that do not affect me.
  - Relevant security advisories fail to reach me.
  - Security advisories are delivered to me shortly after the security vulnerability is detected.
  - I trust all the sources that provide me with security advisories.
  - A central register of security advisories for all platforms is necessary.

B.4 Processing of security advisories

- Asset Management Systems are tools that are used by organizations to efficiently track and maintain their IT inventory. Does your organization use an Asset Management System to keep track of the IT inventory?
  - Yes / No / No, but desired / Unknown
- Do you use automation for processing Security Advisories?
  - Yes / No / No, but desired
- Below is a list of statements about processing security advisories. Please answer to what extent you agree with the following statements:
  (Strongly disagree, Disagree, Undecided, Agree, Strongly agree)
  - Processing of Security Advisories requires organizations to have established vulnerability management procedures.
  - Automation is essential to handle security advisories efficiently.
  - It is easy to match the affected systems of a security advisory to our IT setup.
  - This is an attention test, please mark "Strongly agree".
  - All security advisories should have the same structure.
- Do you forward security advisories to potentially affected end customers and industry partners?
  - Yes / Sometimes / No

B.5 Decision-making based on security advisories

- After reviewing a security advisory, a decision has to be made on how to deal with the described security vulnerability. Are you involved in the decision-making process?
  - Yes / Partially / No
- Below is a list of factors that may influence decision-making regarding the security issues described in an advisory. Please indicate their importance based on your views.
  (Not important, Slightly important, Moderately important, Important, Very important)
  - The likelihood that the vulnerability will be exploited on the company system.
  - The potential impact of the described security vulnerability on the company system.
  - Company policies
  - Available resources
- Which other factors — if any — influence the decision-making process as you see it? (free text)
- Do you use additional sources to reach a decision on how to progress?
  - Yes / Partially / No
- Which additional sources do you use? (free text)

B.6 Additional Remarks

- Are there any additional remarks you would like to add to any of the aforementioned questions? (free text)

B.7 Demographic information

In the last section we would like to capture some demographic information about you, which will enable us to analyze the collected data further.

- What is your year of birth?
- What gender do you identify with?
  - Male / Female / Diverse / Prefer not to say
- What is your country of residence?
- What is your highest education level?
  - Less than high school (no university/college entrance certificate) / Entrance certificate for university/college (high school diploma, GED, GCE, etc.) / Some college, associate degree or equivalent / Bachelor's degree or equivalent / Master's degree or equivalent / Professional degree (M.D., J.D., etc.), Ph.D. (doctoral degree) or equivalent / Prefer not to say / Other: (free text)